

A fast radio burst localized at detection to an edge-on galaxy using very-long-baseline interferometry

Received: 22 July 2023

Accepted: 5 August 2024

Published online: 25 September 2024

 Check for updates

A list of authors and their affiliations appears at the end of the paper

Fast radio bursts (FRBs) are millisecond-duration radio transients whose origins remain unknown. As the vast majority of bursts are one-off events, it is necessary to pinpoint FRBs precisely within their host galaxies at the time of detection. Here we use two purpose-built outrigger telescopes to localize FRB 20210603A at the time of its detection by the Canadian Hydrogen Intensity Mapping Experiment (CHIME). Our very-long-baseline interferometry stations localized the burst to a $0.2'' \times 2''$ final ellipse in the disk of its host galaxy SDSS J004105.82+211331.9. A spatially resolved spectroscopic follow-up revealed recent star formation (H α emission) on kiloparsec scales near the burst position. The excess dispersion measure is consistent with expectations from the nearly edge-on disk of the host galaxy, demonstrating the utility of FRBs as probes of the interstellar medium in distant galaxies. The excess dispersion measure, rotation measure and scattering are consistent with expectations for a pulse travelling from deep within its host galactic plane, strengthening the link between the local environment of FRB 20210603A and the disk of its host galaxy. Finally, this technique demonstrates a way to overcome the trade-off between angular resolution and field of view in FRB instrumentation, paving the way towards plentiful and precise FRB localizations.

Fast radio burst (FRB) 20210603A (Fig. 1) was first detected by the FRB search back end of the Canadian Hydrogen Intensity Mapping Experiment (CHIME)¹, which searches for dispersed single pulses within search beams tiling the 200 deg² primary beam of the telescope. The high signal-to-noise ratio (S/N) of the burst (>100) triggered the recording of voltage data at CHIME² and two small telescopes shadowing a portion of the CHIME field of view (FoV): a single 10 m dish at Algonquin Radio Observatory (ARO10)³ and TONE, a compact array of eight 6 m dishes at Green Bank Observatory⁴. Voltage data dumps from all three stations (Fig. 2) and daily calibration dumps of the Crab pulsar enabled the FRB to be pinpointed to the host galaxy SDSS J004105.82+211331.9 (Fig. 3).

The three stations in our ad hoc array were fixed and shared a common FoV. CHIME is a compact interferometer with 1,024 antennas whose FoV consisted of a $-110^\circ \times -2^\circ$ strip aligned along the local

meridian¹. Like CHIME, ARO10 and TONE drift-scan the sky. They are manually pointed such that their common FoV overlaps the CHIME/FRB search beams at a declination (dec.) of approximately $+22^\circ$ (see also Table 3). Thus, the Crab pulsar can be used as a very-long-baseline interferometry (VLBI) calibrator. When a search beam within the common FoV detects a sufficiently bright single pulse, low-latency alerts trigger dumps of data across the VLBI network ('Instrumentation and observations' in Methods). This observing mode is a technology demonstration for CHIME/FRB outriggers, which will expand the strategy to the full FoV of CHIME using more sensitive outrigger telescopes. Owing to the low sensitivity of ARO10 and TONE, our ad hoc VLBI array has only two useful baselines. Both are largely east–west. Nevertheless, together they provide sufficient (arcsecond-scale) resolution in the north–south direction. The short internal baselines within CHIME and

✉ e-mail: calvin_leung@berkeley.edu




Fig. 1 | The Stokes/dynamic spectrum of FRB 20210603A. We detected the single pulse in autocorrelation at CHIME/FRB with an S/N exceeding 100. The data are shown at a time resolution of $25.6 \mu\text{s}$ with pixel colours scaled to their 1–99 percentile values. To remove dispersion, we used a DM derived by lining up three closely overlapping sub-burst components within the main pulse using fitburst^{19,54}. In addition to the main burst, fainter emission components occurring ~ 12 ms and ~ 18 ms afterwards are visible in CHIME/FRB baseband data but are neglected for VLBI localization. The faint dispersed sweeps left and right of the main pulse are known instrumental artefacts from spectral leakage. The red streaks to the left highlight the frequency channels that are masked out due to radio-frequency interference. Most radio-frequency interference comes from cellular communication and television transmission, bands between 700 and 750 MHz and between 600 and 650 MHz, respectively. a.u., arbitrary units.

TONE do not contain much astrometric information. As they are much shorter than the long interstation baselines, we form voltage beams at CHIME and TONE towards the single pulses of interest. Cable delays for each antenna within CHIME were calculated using the calibration solutions from the CHIME 21 cm back end⁵; for TONE, we used scheduled voltage dumps from the daily transit of Tau A (the Crab nebula) to measure cable delays⁴. Applying these delays allowed the station beams to be formed towards the best-fitting position obtained using CHIME only, which we obtained using the baseband localization pipeline (‘Local calibration and beamforming’ in Methods). We refer to this arcminute-precision position as \hat{n}_0 . After forming station beams, our custom-written VLBI correlator⁶ takes the voltage data from beamformed CHIME, beamformed TONE and ARO10. Within the correlator, geometric delays and Doppler corrections from the Consensus model⁷ were applied to the voltage data in each of the 1,024 frequency channels. We omitted ionospheric and clock corrections from the delay model and calibrated these effects out at the level of visibilities. Our correlator then applied a coherent dedispersion to the Doppler-corrected voltage data from each station. This reduced the effect of intrachannel smearing and narrowed the pulse in time by a factor of a few, which increased the sensitivity of our offline system with respect to the FRB search engine. The VLBI correlation allowed the FRB signal to be pulled out of the noise at the less sensitive stations, where the FRB is undetectable in autocorrelation. After coherent dedispersion and gating, our long-baseline visibilities were generated and written to disk. Despite being undetectable at the outrigger stations in autocorrelation, the FRB was strongly detected with a $S/N \approx 35$ in our visibilities on both the CHIME–ARO10 and CHIME–TONE baselines.

After the burst was detected in cross-correlation, we then applied ionosphere and clock corrections. Typically, these calibration solutions are straightforward to determine using VLBI observations of continuum sources with precisely known positions in the same observing

session. However, with our ad hoc array, such observations are difficult due to the unknown availability of VLBI calibrators at 600 MHz, the fixed pointings and low sensitivity of ARO10 and TONE, and the limited internet connectivity of the ARO10 station. Our calibration strategy, instead, relied on observing bright Crab giant pulses (GPs; Extended Data Fig. 1) once per day, resulting in calibration measurements that are much sparser than typically achievable with a mature, steerable VLBI array. Nevertheless, we conducted monitoring campaigns of the Crab with each baseline individually. We observed and delay-calibrated ten Crab GP datasets on the CHIME–ARO10 baseline and 11 on the CHIME–TONE baseline to empirically estimate our 1σ localization uncertainties (Extended Data Fig. 2). As the Crab emits GPs unpredictably, we observed them in our system with a range of fluences, spectral properties and sky locations. For both the CHIME–ARO10 and the CHIME–TONE monitoring campaigns, the pulses span a range of $\sim 1.1^\circ$ in hour angle. Because our drift-scan telescopes do not track any particular right ascension (RA), the sky rotation and pulse-to-pulse variability mimicked the observation of astrophysical sources with distinct source properties at distinct RAs. The delay uncertainties correspond to a systematic uncertainty ellipse of $0.2 \text{ arcsec} \times 0.2 \text{ arcsec}$ in the east–west and north–south directions respectively (‘VLBI calibration and empirical localization error budget’ in Methods).

In the science run, both the CHIME–ARO10 and the CHIME–TONE baselines operated simultaneously. During this science run, we observed FRB 20210603A and several Crab GPs before and after its detection, which we refer to as C1–C4. These GPs allowed us to derive a set of phase, delay and delay-rate calibration solutions, which we used to localize the FRB (‘FRB localization’ in Methods). However, before performing the localization, we validated the calibration solutions by using them to localize a Crab GP (C3), which we detected 1 day after the FRB and omitted from our calibration solutions, making it an independent check of our calibration (Extended Data Figs. 3–5). The discrepancy between the Crab’s true position and our Crab localization fell well within the systematic uncertainty ellipse from the monitoring campaigns. Finally, we applied the same calibration solutions to localize the FRB (Extended Data Figs. 6 and 7). The target–calibrator separation was 1.5° in hour angle, 0.8° in dec. and 4 h in time. The derived coordinates of FRB 20210603A in the International Celestial Reference System (ICRS) were $\text{RA} \alpha = 0 \text{ h } 41 \text{ min } 05.774 \text{ s} \pm 0.0192 \text{ s}$ and $\text{dec. } \delta = +21^\circ 13' 34.573'' \pm 1.08''$ (Table 1). These coordinates coincide with SDSS J004105.82+211331.9, a disk galaxy with a nearly edge-on orientation (Fig. 3)⁸.

We observed SDSS J004105.82+211331.9 with the Canada–France–Hawaii Telescope (CFHT) MegaCam on 10 September 2021 using the wideband *gri* filter⁹. Figure 3 shows the location of the FRB within the host galaxy. In contrast to other FRB host galaxies that have been robustly identified so far, this galaxy is viewed nearly edge on; it has an inclination of $83 \pm 3^\circ$ (InclinationZoo¹⁰). We determined the *r*-band half-light radius and Galactic extinction-corrected apparent magnitude to be $(8.2 \pm 0.9) \text{ kpc}$ and 17.90 ± 0.01 , respectively, using photometric data provided by the Sloan Digital Sky Survey (SDSS⁸; ‘Host Galaxy Analysis’ in Methods).

Additionally, we acquired spatially resolved spectra with the Gemini multi-object spectrograph¹¹ on 10 August 2021 using a combination of a R400 grating and a GG455 low-pass filter configured with a 1.5 arcsec slit, which covered the wavelength range from 4,650 to 8,900 Å. The slit was co-aligned with the major axis of the galaxy to provide one-dimensional spatial information (Extended Data Fig. 8). Two 1,200 s exposures were taken on the same night but at two different central frequencies, 6,650 and 6,750 Å, to give coverage in the N detector chip gap of the Gemini multi-object spectrograph. The binning was 2×2 , which provided a spatial scale of 0.00292 pix^{-1} and an instrumental resolution of 4.66 Å , sampled at 1.48 Å pix^{-1} . The seeing conditions were very good during the observation night, with a mean airmass of 1.007. Fitting Gaussian line profiles to the H α and N II lines (rest

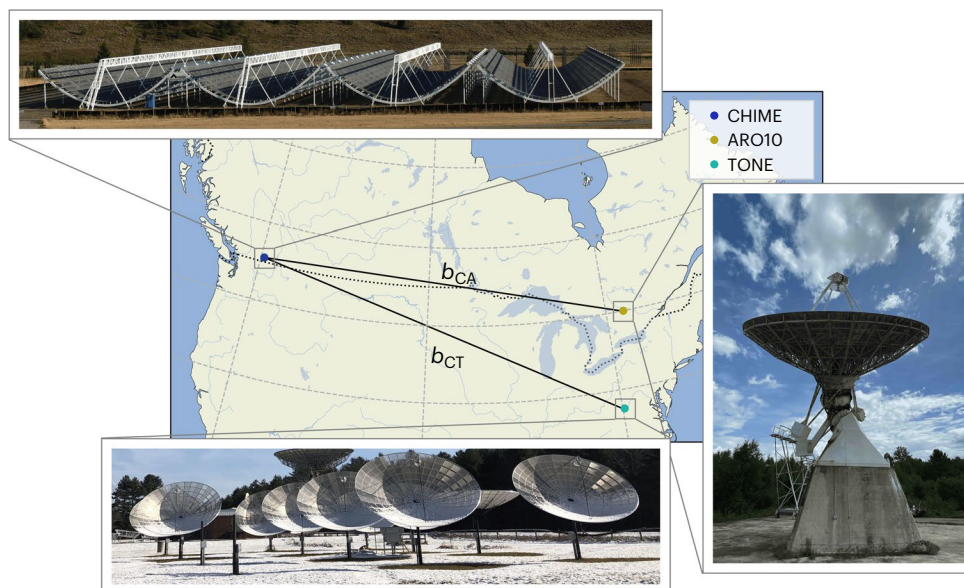


Fig. 2 | Map of baselines formed between CHIME and ARO10 (CA) and TONE (CT). The baselines span from Penticton, British Columbia, to Algonquin, Ontario, and Green Bank, WV, with lengths $b_{CA} = 3,074$ km and $b_{CT} = 3,332$ km.

For our localization analysis, we omit the 848 km baseline between ARO10 and TONE because the FRB was not sufficiently bright to be detected on that baseline. Photographs reproduced with permission.




Fig. 3 | VLBI localization of FRB 20210603A. Left, the 1σ and 2σ localization contours, defined by an empirical estimate of our localization errors using Crab measurements, are overlaid on a CFHT MegaCam *gri*-band image of its host galaxy SDSS J004105.82+211331.9. The nearly edge-on geometry of the host galaxy is apparent. We allow the pixel colours to saturate within ~ 1 half-light radius to accentuate the faint structure on the outskirts of the galaxy. The localization and burst properties point towards a progenitor deep in the ionized disk of the galaxy. Right, $H\alpha$ flux observed at various distances from the galactic

centre along the major axis of the galaxy, calculated from the spectra in Extended Data Fig. 8. Positive (negative) coordinates refer to $H\alpha$ fluxes northwards (southwards) of the galactic centre. Blue circles and upside-down triangles represent detections and 2σ upper limits on the local $H\alpha$ flux, with flux uncertainties estimated using the detrended spectrum (standard deviation for $n = 3,199$). Horizontal bars denote the size of the spectral aperture (1 arcsec). The half-light radius of the galaxy is indicated by a grey shaded area.

wavelengths of 6,564.6 and 6,585.2 Å) yielded a redshift of $z = 0.1772 \pm 0.0001$. Assuming the Planck 2018 cosmology¹², this redshift implies a Galactic extinction- and k -corrected absolute r -band magnitude of -22.03 ± 0.02 .

The spectroscopic redshift of the galaxy (lines shown in Extended Data Fig. 9) implies an angular diameter distance of 639 Mpc and a transverse angular distance scale of $3.1 \text{ kpc arcsec}^{-1}$. Using these values, we measured the projected spatial offset for the FRB of 7.2 kpc

Table 1 | Measured and derived parameters associated with FRB 20210603A and its host galaxy

Parameter	Value
Right ascension α (ICRS)	(10.274058 \pm 0.00008) ^o
Declination δ (ICRS)	(21.226270 \pm 0.0003) ^o
CHIME arrival time at 400 MHz	15:51:34.431652 UTC on 3 June 2021
Dispersion measure (DM)	(500.147 \pm 0.004) pc cm ⁻³
DM _{MW-NE2001} [†]	(40 \pm 8) pc cm ⁻³
DM _{MW-halo} [†]	(30 \pm 20) pc cm ⁻³
DM _{cosmic}	(172 \pm 90) pc cm ⁻³
(DM _{host})/(1+z) = (DM _{host-disk} + DM _{host-halo})/(1+z)	(257 \pm 93) pc cm ⁻³
Rotation measure (RM)	(-219.00 \pm 0.01) rad m ⁻²
RM _{MW} [†]	(-22.4 \pm 0.3) rad m ⁻²
RM _{iono} [†]	+1.4 rad m ⁻²
$\Pi_{L-800\text{MHz}}$	$\geq 96\%$
$\Pi_{L-400\text{MHz}}$	$\geq 87\%$
$\tau_{600\text{MHz}}$	(165 \pm 3) μ s
$\tau_{600\text{MHz-NE2001}}^{\dagger}$	1.02 μ s
Fluence	(64.4 \pm 6.5) Jy ms
Flux density	(64.9 \pm 6.5) Jy
Specific energy	5.7 $\times 10^{31}$ erg Hz ⁻¹
Specific luminosity	5.8 $\times 10^{34}$ erg s ⁻¹ Hz ⁻¹
Band-averaged pulse FWHM	740 μ s
Spectroscopic redshift, z	0.1772 \pm 0.0001
Photometric redshift, z _{phot} [†]	0.175 \pm 0.0133
Inclination angle	(83 \pm 3) ^o
Present-day stellar mass, log(M _{host} [*] /M _o)	10.93 ^{+0.04} _{-0.04}
Metallicity, log(Z/Z _o)	-0.22 ^{+0.05} _{-0.04}
Mass-weighted age	4.32 ^{+0.73} _{-0.75} Gyr
Total star formation rate (SFR)	$\geq 0.24 \pm 0.06 M_{\odot} \text{ yr}^{-1}$
Projected offset	7.2 kpc
r-band half-light radius	(8.2 \pm 0.9) kpc
Absolute r-band magnitude	-22.03 \pm 0.02
E(B-V)	0.28

Properties derived from radio and optical follow-up data are listed in the top and bottom halves of the table, respectively. Parameters derived from external models or measurements are indicated with daggers. z_{phot}, DM, τ and RM_{iono} predictions are from ^{20,21,56,60,77,78}.

from the host galactic centre along the host galactic plane. This offset is consistent with the distribution of projected offsets measured from a sample of both repeating and non-repeating FRBs localized by the Australian Square Kilometre Array Pathfinder (see, for example, Fig. 9 in ref. 13), with the caveat that our localization ellipse is too large to draw any meaningful conclusion about the host offset. To characterize the host galaxy, we combined the Gemini spectra with archival photometry from the Two Micron All Sky Survey (2MASS)¹⁴ and the Wide-Field Infrared Space Explorer (WISE)¹⁵ to extend our wavelength coverage upwards to $1 \times 10^5 \text{ \AA}$ ('Host galaxy analysis' in Methods).

We fitted a spectral-energy distribution model to the combined spectral and photometric data using the Bayesian fitting package Prospector¹⁶. We estimated best-fitting values and uncertainties for

Table 2 | Priors for modelling the spectral-energy distribution

Parameter		Prior [min, max]
log(M [*] /M _o)	Present-day stellar mass	log uniform [10, 12]
log(Z/Z _o)	Metallicity	Top hat [-2, 0.19]
t	Time since formation (Gyr)	Top hat [0.1, 12]
τ	Star formation characteristic decay rate (Gyr)	log uniform [0.3, 15]
dust2	Diffuse dust V-band optical depth	Top hat [0, 3]

These parameters were used to model the host galaxy with a delayed- τ model as implemented in Prospector.

the present-day stellar mass, mass-weighted age, V-band dust extinction and metallicity of our host galaxy using Markov-chain Monte Carlo posterior sampling (Extended Data Fig. 9)¹⁷. The parameters determined by Prospector and the star formation rate (SFR) are shown in Table 2. From the H α luminosity measured with Gemini data, we determined the galaxy's overall SFR ($0.24 \pm 0.06 M_{\odot} \text{ yr}^{-1}$) and detected star formation in the ~ 10 kpc-scale vicinity of the FRB. The detection of H α emission is potentially a sign of recent star formation (~ 10 Myr) and young stellar populations. However, as for other FRBs, spatially resolved spectroscopic studies of this galaxy are needed to further constrain the age and nature of the FRB progenitor.

In addition to the host galaxy properties, the burst itself can provide insights into the sight line towards the FRB progenitor and the progenitor itself. For instance, if the FRB were in the inner disk, it would experience enhanced dispersion and scattering due to the long line-of-sight path out of the host galaxy's ionized disk towards the observer, like pulsars at low Galactic latitudes in the Milky Way. FRB 20210603A, therefore, allows for a detailed accounting of host galactic contributions to the observed dispersion measure (DM), rotation measure (RM) and scattering (pulse broadening). To check this possibility, we calculated the DM excess by subtracting estimated DM contributions from the Milky Way, the Milky Way halo and the intergalactic medium from the measured DM. We obtained a large DM excess of DM_{host}^r = (302 \pm 109) pc cm⁻³, where the superscript denotes that DM_{host}^r is defined in the host galaxy's rest frame.

One interpretation for this excess would be a dense environment local to the FRB progenitor¹⁸, which may add measurable contributions to the DM, RM or scattering timescale. Another interpretation is that the host galaxy itself dominates the DM excess, with subdominant circumburst contributions to the other properties. Our estimate of the DM budget of the host galaxy is $\sim (264 \pm 97) \text{ pc cm}^{-3}$ ('Dispersion and scattering analysis' in Methods), which is consistent with the latter hypothesis. Although both interpretations are compatible with the data in hand, Occam's razor leads us to favour the interpretation that the excess DM of this FRB is dominated by the host galaxy's disk (Extended Data Fig. 10).

This is consistent with our measurements of the pulse broadening timescale, which we determined by fitting a pulse model to the FRB's dynamic spectrum. The complex time-frequency structure of the bright main burst required three subpulse components, temporally broadened by the same characteristic timescale, to obtain a robust fit to the data ('Burst morphology' in Methods and ref. 19). This placed an upper limit on the scattering timescale of $\tau_{600\text{MHz}} \lesssim 165 \pm 3 \text{ \mu s}$ at a reference frequency of 600 MHz. As the scattering from the Milky Way is expected to be subdominant ($\sim (1.0 \pm 0.5) \text{ \mu s}$)^{20,21}, we concluded that the observed pulse broadening was dominated by an unresolved substructure in the burst profile or extragalactic scattering, probably in the host rather than the Milky Way²². If the measured broadening timescale is attributed entirely to scattering and scaled to the rest frame and scattering geometry of the host galaxy, the implied scattering efficiency of the host galactic gas is like that of a typical sight line towards a

pulsar through a galactic disk with Milky Way-like density fluctuations ('Dispersion and scattering analysis' in Methods).

In addition, the interpretation of a dominant host galactic contribution is consistent with our measurement of the burst RM ('Polarization analysis' in Methods). After subtracting Galactic and terrestrial contributions (RM_{MW} and RM_{iono} in Table 1), the excess $RM_{excess} = (+198 \pm 3) \text{ rad m}^{-2}$. As no intervening systems (for example, galaxy groups or clusters) have yet been observed along this sight line, the RM contribution from the intergalactic medium is probably negligible²³. The magnitude of the RM excess is unremarkable and can easily be explained by contributions from the host galaxy's interstellar medium. These properties suggest that the source of FRB 20210603A is close to its galactic plane (Extended Data Fig. 10), consistent with our localization ellipse.

In conclusion, we have commissioned a VLBI array to demonstrate the first VLBI localization of a non-repeating FRB. The limitations of our ad hoc VLBI array, however, leads to a final localization uncertainty on par with connected-element interferometers like the Australian Square Kilometre Array Pathfinder, DSA-110 and MeerKAT. Nevertheless, this paves the way towards precisely localizing a large sample of one-off bursts using VLBI. The FRB 20210603A sight line has implications for galactic astrophysics and the progenitors of FRBs. It demonstrates the potential for using edge-on FRB host galaxies as probes of the ionized gas of other galaxies. In addition, the H α emission in the neighbourhood of the FRB suggests recent star formation activity. This highlights the need for a high-resolution follow-up to discriminate among progenitor models by assessing whether FRBs are spatially coincident with star-forming regions²⁴. The instruments and methods used here constitute pathfinders for the CHIME/FRB outtriggers project, which will enable VLBI localizations of large numbers of both repeating and non-repeating sources^{3,25,26}. Thus, a more complete picture of the diverse host environments of FRBs, and how the environments correlate with other burst properties, will soon be available.

Methods

Instrumentation and observations

We used a VLBI network consisting of three stations: CHIME at the Dominion Radio Astrophysical Observatory (DRAO)¹, ARO10 (a single 10 m dish at Algonquin Radio Observatory³) and TONE (a compact array of eight 6 m dishes at Green Bank Observatory⁴). CHIME/FRB detected FRB 20210603A at 15:51 UTC on 3 June 2021. Figure 1 shows the Stokes / dynamic spectrum of the beamformed data from FRB 20210603A as observed at CHIME. Between August 2018 and May 2021, 35.6 h of exposure were accumulated in the direction of FRB 20210603A; however, only the burst reported here was detected. To calibrate the VLBI calibration and test our localization procedure, we used several Crab GPs captured at a cadence of one per day, which we refer to as C1–C4 (Extended Data Fig. 1).

CHIME/FRB. CHIME consists of four 20 m \times 100 m cylindrical paraboloid reflectors oriented with the cylinder axis in the north–south direction². Each cylinder is fitted with 256 dual-linear-polarization antennas that are sensitive in the frequency range 400–800 MHz. The 2,048 analogue signals from the antennas are amplified and digitized using an array of 128 motherboards with field-programmable gate arrays with mezzanine analogue-to-digital converters called ICE boards²⁷. At each ICE board, raw voltages are channelized with a polyphase filterbank producing 1,024 complex channels with 2.56 μ s time resolution. We refer to the channelized and time-tagged voltage data as raw baseband data (as opposed to beamformed baseband data; 'Local calibration and beamforming' in Methods). These data are sent to 256 GPU-based compute nodes comprising the X-Engine correlator driven by the open-source Kotekan software repository^{28,29}. We computed the spatial correlation and summed the polarizations to form 1,024 (256 N–S \times 4 E–W) independent beams within the north–south primary

Table 3 | A summary of the properties of the CHIME, ARO10 and TONE stations

Property	CHIME	ARO10	TONE
SEFDs S_{sys} at $\delta = +22^\circ$	–	$\sim 1.7 \text{ kJy}$	$\sim 20\text{--}40 \text{ kJy}$
Field of view (at 600 MHz)	$\sim 110^\circ \text{ N--S,}$ 1.74° E--W	3.59°	$\sim 6\text{--}11^\circ$
Processed frequency channels	916	1,024	1,024
Baseline length	–	$b_{CA} = 3,074 \text{ km}$	$b_{CT} = 3,332 \text{ km}$
Longitude (deg)	-119.6237	-78.0701	-79.8452
Latitude (deg)	49.3207	45.9556	38.4293

The system equivalent flux density (SEFD) at ARO10 was calculated with a set of Crab GPs³. The SEFD and FoV of TONE were computed from a drift-scan observation of Taurus A (ref. 4). The CHIME SEFD at $+22^\circ$ has not been calculated, but its system temperature has been extensively studied in ref. 2.

beam³⁰. These beams were searched for FRBs in real time using detection pipelines designed to discover radio transients. The real-time pipeline and the baseband system collectively make up the CHIME/FRB instrument^{1,5}. The baseband system uses a memory ring buffer system to record (or 'dump') baseband data to disk. The ring buffer holds $\sim 35.5 \text{ s}$ of baseband data for subsequent capture by a detection trigger. On successful detection of an FRB candidate by the real-time pipeline above an S/N of 12, a trigger from the real-time pipeline saved a $\sim 100 \text{ ms}$ snapshot of data centred around the pulse at each frequency channel of the baseband buffer. The latency between the time of arrival of a signal and the triggered baseband recording was typically $\sim 14 \text{ s}$. The buffer can record the full band's worth of data when the dispersive sweep of the FRB does not exceed $\sim 20 \text{ s}$ (corresponding to a maximum DM of $\sim 1,000 \text{ pc cm}^{-3}$).

Upon detection by the real-time pipeline of an FRB or a Crab pulsar GP³¹ in the FoV of TONE and ARO10, a trigger was sent to the active outtriggers. To prevent GP triggers overwhelming the baseband read-out system with thousands of events, we recorded only triggers with a detection S/N greater than 40 (near CHIME's zenith) and having a duty cycle of 1%. This resulted in a Crab GP dump rate of about once per day at each station. We have summarized in Table 3 the stations and their properties, and now discuss the two outtrigger stations in detail.

Algonquin Radio Observatory 10 m telescope. ARO10, a single 10 m dish, is at the Algonquin Radio Observatory in Algonquin Provincial Park, Ontario. The CHIME–ARO10 baseline $b_{CA} \geq 3,000 \text{ km}$ (Fig. 2). The two analogue signals from the polarizations of the single CHIME cloverleaf feed³² were digitized and acquired with a digital infrastructure identical to that of CHIME and TONE except that the large ($\sim 24 \text{ h}$ long) ring buffer was stored on hard disks. A complete description of the radio-frequency chain and the digital system is provided elsewhere³. The data from ARO10 exhibit a delay drift relative to DRAO amounting to $\sim 0.1 \mu\text{s}$ per day. This extra shift in addition to the $\sim 2 \text{ ms}$ geometrical delay is predictable and is corrected for (Fig. 15 of ref. 3).

TONE. TONE is at Green Bank Observatory near the Green Bank Interferometer Control Building. The CHIME–TONE baseline is $b_{CT} \approx 3,332 \text{ km}$ long (Fig. 2). TONE is an array of 6 m dishes placed in a regular 4×3 grid with 9.1 m spacing with the shorter side aligned 60° off true north. Each dish was oriented to observe the Crab pulsar at the same time as CHIME. Eight dishes were deployed with feeds instrumented with active-balun dual-polarized cloverleaf antennas^{32,33}. The 16 analogue signals were each transmitted over a radio-frequency-over-fibre system³⁴. For this work, seven signals from one polarization and six signals from the other were used to synthesize a single beam for VLBI. The signals from the radio-frequency-over-fibre receiver were digitized and channelized by an ICE board (in the same way as previously described for CHIME and ARO10). A TM-4 GPS clock module³⁵ provided a 10 MHz clock and

absolute time. Additionally, a 10 MHz maser signal was fed into the ICE board²⁶ to replace one of the analogue inputs for post hoc clock delay characterization³⁶. The digitized and channelized voltages were sent through two 40 Gbit ethernet network links over to the recording computer node. The recording node used Koteikan, as do those at CHIME and ARO10, to create an ~40 s buffer of baseband data²⁵. The length of the buffer must accommodate both the latency of the CHIME/FRB detection pipeline and the network in addition to the science data. The baseband read-out saved an ~0.5 s slice of the buffer around the pulse on the arrival of a trigger to disk for offline VLBI analysis. Taurus A was used as a calibrator to phase the antennas within TONE for beamforming ('Local calibration and beamforming' in Methods). See ref. 37 for a detailed description of the system and its performance.

Clock calibration

Timing errors are intrinsic to the digital back ends at each station, which are locked to different clocks with varying degrees of stability. The severity of the timing errors depends on the type of clock used at each station and varies from unit to unit. Timing errors are characterized in terms of the Allan deviation ($\sigma(\Delta t)$) as a function of timescale Δt , for example, between successive clock calibrations²⁶. The CHIME digital system was locked to a single 10 MHz clock signal provided by a GPS-disciplined, oven-controlled, crystal oscillator. Although sufficient for the operations of CHIME as a stand-alone telescope, this clock does not meet the stringent stability requirements for VLBI with CHIME/FRB outriggers. To overcome this limitation, we sampled the more stable passive hydrogen maser (at DRAO) during FRB VLBI observations²⁶ on a regular cadence. This minimally invasive clocking system was developed as part of the effort to expand CHIME's capabilities to include VLBI with CHIME/FRB outriggers. It works by digitizing the signal from an external maser using one of the inputs of the GPS-clock-driven ICE board. We read out a 2.56 μ s snapshot of maser data at a cadence of once every $\Delta t_{\text{GPS,C}} = 30$ s at CHIME. The data read out from the maser were processed offline to measure the drift of the GPS clock between calibrator observations. A similar read-out system recorded a 10 MHz clock at TONE at a cadence of $\Delta t_{\text{GPS,T}} = 1$ s. In contrast, the digital system of ARO10 was directly clocked by an actively stabilized hydrogen maser, which removed the need for station-based clock corrections.

Once clock corrections are applied to the observations, the expected delay error between two observations separated by Δt_{sep} in time is given by the quadrature sum of the jitter at each station. Assuming that the jitter is characterized by the Allan deviation of the maser alone, this is given by $\sigma_{\text{maser}}(\Delta t_{\text{sep}})\Delta t_{\text{sep}}$. On 24 h timescales, this corresponds to a delay error of ~0.35 ns for the CHIME–ARO10 baseline (one passive and one active maser) and ~0.48 ns for the CHIME–TONE baseline (two passive masers)²⁶. In addition, on the CHIME–TONE baseline, observations are referenced to the maser by interpolating between the maser read-outs directly before and after the observation. The slow cadence of the maser read-out at these stations induced a further interpolation error of size $\sigma_{\text{GPS}}(\Delta t_{\text{sync}}) \times \Delta t_{\text{sync}}$ (ref. 36), for a total of 0.52 ns.

Local calibration and beamforming

CHIME has 1,024 antennas, and TONE has eight antennas. It is infeasible to correlate such a large number of antennas as independent VLBI stations. To reduce the computational burden of correlating such a large array, we coherently added, or beamformed, the raw baseband data from the antennas within each station to combine the low-sensitivity antennas from a single station into a high-sensitivity equivalent single dish using beamforming.

Beamforming requires independent measurements of the individual sensitivities and delays for each antenna, that is, complex-valued gains that contain both amplitude and phase information. At CHIME, the infrastructure to calculate these so-called N^2 -gains and a tied-array beamformer have already been developed². We generalized several of

CHIME's software frameworks^{38–41} to use the same basic N^2 -gain calibration algorithms⁴² at TONE. First, the visibility matrix from all N^2 pairs of antennas at the correlator is calculated when a bright point source (Taurus A for TONE) dominates the FoV. In the single-source limit, the visibility matrix has a rank-1 eigendecomposition; the non-singular eigenvector and eigenvalue encode a combination of geometric delays and instrumental gains and delays. Once the gains are characterized, they are used to beamform the raw baseband data from CHIME and TONE towards the best-known positions of the Crab and the FRB provided by the baseband pipeline (\hat{n}_0). The synthesized beam at CHIME is ~1 arcmin wide, and the synthesized beam at TONE is ~0.5° wide. As the FRB's true position is well within a synthesized beam-width away from \hat{n}_0 , our final sensitivity depended only weakly on \hat{n}_0 .

VLBI correlation

After beamforming was complete at each station, the beamformed baseband data were correlated with a custom Python-based VLBI correlator⁶. We used the stand-alone delay model implemented in difxcalc⁴³ to calculate geometric delays towards the fiducial sky location \hat{n}_0 of each source. For the Crab pulses, we used the VLBI position of the Crab pulsar⁴⁴ extrapolated using its proper motion to the epoch of our observations:

$$\hat{n}_0 = (83.6330379^\circ, 22.014501^\circ), \quad (1)$$

with the RA and dec. reported in decimal degrees. Including the pulsar position error ($\sigma_{\hat{n}}$) and the proper motion (μ) error (σ_{μ}) extrapolated over ~10 yr from recent Crab pulsar astrometry⁴⁴, we summed the absolute position error at the archival observing epoch and the uncertainty in the proper motion, scaled by the time between our observations (~10 yr), in quadrature for the RA and dec. individually. The uncertainties in the Crab position propagated into equally sized positional uncertainties of the FRB; however, these are subdominant compared to our systematics, so we did not quote them above. For the FRB, we used the best-fitting position derived from a CHIME-only baseband localization ($\hat{n}_0 = (10.2717^\circ, 21.226^\circ)$). This is precise to within an arc-minute; nevertheless, we found strong fringes on the FRB pointing towards this position.

In our correlator, we broke the total delay into an integer number of 2.56 μ s frames and a subframe (or subinteger) component whose value is in the range ~1.28 to ~1.28 μ s. The integer shift was applied to the data through an array shift, and the subinteger shift was applied by a phase rotation to each 2.56 μ s frame. Although this time resolution is lower than that of more conventional VLBI back ends, compensating for the delay on this timescale did not appreciably increase the phase errors, even at the top of the band where these would be most noticeable. We estimated an upper limit on the phase error at the top of our band to be $\sim \epsilon \times 2.56 \mu\text{s} \times 800 \text{ MHz}$, where ϵ is the maximum delay rate encountered during our observations. For the most extreme scenario of two antipodal VLBI stations at the equator, $\epsilon \approx 3 \times 10^{-6}$ gives a phase error of 2.2°, an acceptably small amount of decorrelation.

After delay compensation, each of the 1,024 frequency channels of data was de-smeared by a coherent dedispersion kernel⁴⁵. Although several conventions could have been used (see, for example, equation (5.17) in ref. 46), we used the following kernel in our VLBI correlator:

$$H(\nu) = \exp\left(2\pi i k_{\text{DM}} \text{DM} \frac{\nu^2}{2\nu_k^2(\nu_k + \nu)}\right). \quad (2)$$

In equation (2), we took $k_{\text{DM}} = 1/(2.41 \times 10^{-4}) \text{ s MHz}^2 \text{ pc}^{-1} \text{ cm}^3$ (for consistency with previous conventions in the pulsar community^{46,47}). The fiducial DM of the FRB was taken to be $(500.147 \pm 0.004) \text{ pc cm}^{-3}$. We chose this dedispersion kernel to avoid introducing delays into each frequency channel (as it preserves times of arrival at the central frequency of each channel). The chosen DM de-smeared the pulse

within each frequency channel, which concentrated the signal into a narrow temporal duration and increased the correlation power. The argument $v \in [-195.3125 \text{ kHz}, +195.3125 \text{ kHz}]$ indicates the offset from the reference v_k , chosen to be the centre of each frequency channel: $v_k \in [800.0, 799.609375, \dots, 400.390625] \text{ MHz}$.

After the delay compensation towards the fiducial sky position $\hat{\mathbf{n}}_0 = (\alpha_0, \delta_0)$ and coherent dedispersion, we formed visibilities for each frequency channel (indexed by k) independently on both long baselines involving CHIME (b_{CA} and b_{CT} , hereafter indexed by i) by multiplying and integrating the complex baseband data. To reject noise, we integrated only $\sim 100 \mu\text{s}$ of data on either side of the pulse in each of the 1,024 frequency channels. In addition, we rejected channels with radio-frequency interference (Extended Data Fig. 1) within each site. This produced ~ 900 complex visibilities per baseline, which were used for localization (hereafter referred to as $V[i, k]$). We integrated 13 other windows of the same duration in the same dataset but shifted to off-pulse times to estimate the statistical uncertainties on the visibilities. The statistical uncertainties are hereafter referred to as $\sigma[i, k]$.

VLBI calibration and empirical localization error budget

The complex visibilities $V[i, k]$ must be phase-calibrated before the localization analysis. We calibrated the visibilities with phase, delay and rate corrections derived from our Crab GPs before performing our final localization analysis. In an ideal set-up, we would systematically characterize localization errors in the CHIME–ARO10–TONE array as a function of sky pointing and time separation and perform end-to-end localization of known pulsars as a check of our localization. However, our ability to do so is limited due to logistical factors at each station. Perhaps most logistically difficult is the extremely limited internet access at the ARO10 site, which fundamentally limited the amount of data that could practically be read out from the ARO10 site³. At TONE, frequent misalignments of the dishes due to high winds requires manual repointing and recalibration of the array, which frequently interrupts observations. Therefore, the only data available for characterizing the full CHIME–ARO10–TONE array around the time that the FRB was observed are a sequence of triggered baseband dumps from the Crab pulsar collected in May–June 2021, simultaneous with CHIME, which occurred at a cadence of about one per day, at each station. We enumerated these Crab pulses as C1–C4. Waterfall plots of these pulses, in addition to the FRB, are shown in Extended Data Fig. 1.

Within the constraints of these limited data, we performed the following steps for VLBI calibration. We used C2, the closest Crab pulse in time to the FRB, as a delay and phase calibrator. Thus, we calculated the instrumental phase and delay solutions for all baselines and applied them to all observations on all baselines. The phase and delay solutions removed static instrumental cable delays and frequency-dependent beam phases and suppressed unwanted astrometric shifts related to baseline offsets towards the elevation angle of the Crab, which is less than a degree away from the FRB in alt-azimuth coordinates. In addition to the phase and delay calibration, a large delay-rate correction ($\sim 0.1 \mu\text{s}$ per day) is needed for the CHIME–ARO10 baseline³. Upon removal of the CHIME–ARO10 clock rate, our delay residuals were small (Extended Data Fig. 2). That figure also includes all the delay residuals from historical data available on each baseline individually, which were calibrated similarly (with a clock rate correction for CHIME–ARO10 and with no clock rate correction detected for CHIME–TONE).

In the absence of commissioning data available when all three stations were operating, we characterized each baseline individually. For CHIME–ARO10, we used a previously published dataset of ten correlated Crab pulses from October 2020. For CHIME–TONE data, we used 11 Crab GPs from February and March 2021 when the instrument was commissioned⁴. From these data, we established the 1σ systematic localization uncertainties by calculating the root-mean-square (r.m.s.) delay errors on each baseline using most of the data plotted in Extended Data Fig. 2. The r.m.s. delay errors for the CHIME–ARO10 and CHIME–TONE

baselines were 8.5 and 6.0 ns, respectively. These errors were calculated from ten and 11 Crab single-baseline measurements, respectively. They were calculated by excluding the pulses used for the delay and rate calibration (whose delay residuals were zero by definition) and faint pulses (CHIME–TONE data from March 2021) whose fringe detections were marginal due to a windstorm at Green Bank, which blew several TONE dishes off-axis before they were manually reported.

Crab localization

In addition to quantifying the delay errors on each baseline individually using Crab pulses, we performed an independent, end-to-end cross-check of the delay and rate solutions derived for the FRB using C3. This was the only remaining Crab GP that had been detected at all stations and baselines as it was not used to obtain delay and rate solutions. We used it as an independent check of our delay and rate solutions and of our localization procedure, which combines data from both baselines.

To localize C3, we calibrated C3 visibilities for both baselines using the aforementioned delay and phase solutions from C2. In addition, for the CHIME–ARO10 baseline, we applied the clock rate measured from C1 and C2. The calibrated residual delays when the C3 data were correlated towards the true Crab position were 2.8 ns for the CHIME–ARO10 baseline and 2.1 ns for the CHIME–TONE baseline. To further model the short-term trend seen in the CHIME–TONE delay residuals, we attempted to apply a clock rate correction to CHIME–TONE data measured from C2 and C4 (as the TONE correlator restarted between C1 and C2). Doing so changed the CHIME–TONE delay by only ~ 1 ns. The residual delays, as well as the final delay-rate correction, were subdominant to our 1σ systematic error budgets of 8.5 and 6.0 ns for the CHIME–ARO10 and CHIME–TONE baselines.

We refer to the visibilities calibrated this way as $v[i, k]$ (not to be confused with the uncalibrated visibilities $V[i, k]$), where i denotes the baseline (either CA or CT) and k denotes our 1,024 independent frequency channels. They are plotted with residual delays removed in Extended Data Fig. 3. In addition to the correlation start times in each channel $t_0[i, k]$ and the baseline vectors \mathbf{b}_{CA} and \mathbf{b}_{CT} , we used $v[i, k]$ to localize C3 to an inferred position $\hat{\mathbf{n}}$ relative to the fiducial sky position ($\hat{\mathbf{n}}_0$) used to correlate C3.

Several approaches to localizing single pulses have been taken in the literature^{3,25,48}, which reflects the nontrivial challenge of astrometry with sparse uv coverage. For example, the traditional method of making a dirty map of a small field and using traditional aperture synthesis algorithms to deconvolve the point spread function is not suited well to the present VLBI network because of the sparse uv coverage. We found that one robust method is to take the delay estimated from the peak of the Fourier transform of the visibilities and use that delay measurement to localize the FRB by maximizing equation (3). This method is robust in the sense that equation (3) has only one global maximum, so it works well even when the true position is at minutes away from the initial pointing.

$$\log \mathcal{L}_\tau = \sum_{i=\text{CA,CT}} \frac{(\tau_i^{\text{max}} - \tau_i(\hat{\mathbf{n}}))^2}{2\sigma_{\tau,i}^2} \quad (3)$$

The drawback of this simple method is that it is sensitive to information in only the linear part of the phase model ($d\phi/dv_k$), which means that it mixes the ionospheric and geometric delays and, therefore, is accurate only at the arcsecond level. Working in visibility space is a straightforward way to break this degeneracy, as we can fit higher-order contributions to the phase as a function of frequency. We fitted equation (4) to our data to disentangle the ionospheric from the geometric delays:

$$\phi[i, k] = 2\pi \left(v_k \tau_i + k_{\text{DM}} \Delta \text{DM}_i \frac{1}{v_k} \right). \quad (4)$$

We obtained the best-fitting solution by maximizing the visibility-space likelihood function (equation (6)). Practically, it was difficult to do this because the posterior was highly multimodal, as seen in our final contours in Extended Data Fig. 4. We resorted to using a box centred on a good initial guess. For the RA and dec., the initial guess was taken from the \mathcal{L}_τ localization. The initial guesses for ΔDM for each baseline were determined by independently optimizing the S/N (equation (5)) over a range of ΔDM and delay values for each baseline:

$$\rho_{\text{sf}}(\tau, \Delta\text{DM}) = \left\| \sum_k \frac{\mathcal{V}[i, k] \exp(-i\phi[i, k])}{\sigma[i, k]} \right\|. \quad (5)$$

With these initial guesses, we evaluated equation (6) on a four-dimensional grid to simultaneously solve for the source position and the ionosphere parameters. Equation (6) uses a signal-to-noise weighting scheme, weighting the real part of the phase-rotated visibilities by $|\mathcal{V}|/\sigma^2$. The denominator of this weighting corresponds to inverse noise weighting, and $\sigma[i, k]$ refers to the statistical uncertainties in the visibilities. The numerator corresponds to an upweighting by the visibility amplitude. As the FRB was detected in each channel with an S/N of ~ 5 – 10 and as it was the single dominant source of correlated flux in the correlated data, we used the visibility amplitude $|\mathcal{V}[i, k]|$ as a convenient approximation to the statistically optimal upweighting, which is the true signal power in each channel after applying appropriate band-pass and beam corrections to each baseline. Note that the band-integrated S/N reported elsewhere (for example, see Extended Data Fig. 3) is an underestimate of the true S/N, as the flux from the FRB contributed measurably to the r.m.s. noise level of the fast Fourier transform.

$$\log \mathcal{L}_\phi \propto \sum_{i=\text{CA, CT}} \sum_{k=0}^{1,023} \frac{\|\mathcal{V}[i, k]\| \text{Re}[\mathcal{V}[i, k] \exp(-i\phi[i, k])]}{\sigma[i, k]^2}. \quad (6)$$

The posterior as a function of our four parameters (α , δ , $\Delta\text{DM}_{\text{CA}}$ and $\Delta\text{DM}_{\text{CT}}$) is shown in Extended Data Fig. 4. We took the parameter set that maximizes the likelihood on the grid as the best-fitting model. The model phases corresponding to these parameters, as well as the model phases corresponding to the parameters that maximize \mathcal{L}_τ , are plotted in Extended Data Fig. 3. The maximum \mathcal{L}_ϕ position of C3 is $\hat{\mathbf{n}} = (83.633053^\circ, 22.014539^\circ)$. Finally, we draw systematic error contours around this best-fitting position using $\sigma_{\tau, i} = 8.6$ and 6.0 ns in Extended Data Fig. 5. The 1σ systematic error contour drawn around the best-fitting position easily encloses the Crab's true position and the delay-only best-fitting position, but does not separate out the ionospheric delay, showing that the ionosphere is not the dominant source of systematic error in our localization.

FRB localization

We applied the exact same calibration solutions used to localize C3 to the FRB visibilities. Following the same procedure, we used the coarse localization with \mathcal{L}_τ to coarsely localize the FRB. The \mathcal{L}_τ position was $\hat{\mathbf{n}} = (10.274056^\circ, 21.22624^\circ)$ which is offset from the baseband localization by 8 arcsec in the RA direction and approximately -1.3 arcsec in the declination direction. To recover some sensitivity, we repositioned the correlator phase centre towards this refined position before fitting the fringes of the calibrated visibilities (Extended Data Fig. 6) for the ionosphere using \mathcal{L}_ϕ . The initial guesses for the ionosphere were estimated as done previously for C3. The fringe fit yielded the maximum-likelihood position $\hat{\mathbf{n}} = ((10.274058 \pm 0.000080)^\circ, (21.226270 \pm 0.000300)^\circ)$ (Table 1). The posteriors are shown in Extended Data Fig. 7.

Possible error sources. Next, we summarize some known contributions to our systematic error, as we could not account for the empirically measured delay errors (8.5 and 6.0 ns at 1σ). We have seen that

these correspond to a $0.2'' \times 2''$ ellipse on the sky, and that relative to this ellipse, the effect of including the ionosphere was small. We estimated our station-positioning errors to be 21 mas by assuming a conservative ~ 10 m baseline uncertainty. Time variations in the phasing of the antennas may also have occurred at CHIME or TONE, as the relative cable lengths fluctuate on weeklong timescales by around 0.1 ns at these stations, although they are recalibrated every day. Uncertainties in the proper-motion extrapolated position of the Crab pulsar at its current epoch (2 mas) were also subdominant. Another systematic uncertainty was the astrometric frame tie between our VLBI localization (ICRS) and optical follow-up observations, which were performed relative to the FK5/ICRS reference frame. The discrepancy between the frames was of the order of -1 mas (refs. 49–52).

As none of these explain the magnitude of our systematic error, we have to consider alternate sources of delay fluctuations. One-day timescale variations in the masers' relative oscillation frequencies or the signal chains carrying the maser signals to the digitizers in the F-engine may have added delay noise on timescales relevant for our sparse calibration. Regardless, our empirical measurement of the r.m.s. delay residuals (Extended Data Fig. 2), which were used to quantify our localization uncertainty, encompasses all the known and unknown sources of systematic astrometric uncertainties, putting our scientific conclusions on a firm footing. In the future, dedicated lab tests could verify this. To avoid the issue completely, the time between VLBI calibrations could also be shortened to minutes or hours. With future outrigger stations having an order of magnitude collecting area than ARO10 and TONE, this will be readily achievable.

Burst morphology

FRB 20210603A was detected with a S/N of ~ 136 in the CHIME/FRB real-time detection pipeline. Afterwards, we characterized its burst morphology and estimated its brightness using high-resolution baseband data. The flux, fluence, specific energy and specific luminosity of the burst are listed in Table 1. Viewed in baseband data, the FRB has a broadband main pulse with a total full-width at half-maximum (FWHM) of 740 μs . In addition, two trailing components are visible in the baseband dump (Fig. 1). Using the DM_phase algorithm⁵³, we lined up substructures in the main pulse, yielding a DM of (500.147 ± 0.004) pc cm⁻³. The DM and the baseband data were input to fitburst⁵⁴, which simultaneously fitted the main burst with three closely spaced sub-bursts with FWHM widths of 310 , 450 and 834 μs , all broadened by 165 μs at 600 MHz.

Dispersion and scattering analysis

In general, the observed DM of an FRB can be split into four components:

$$\text{DM}_{\text{FRB}} = \text{DM}_{\text{MW-disk}} + \text{DM}_{\text{MW-halo}} + \text{DM}_{\text{cosmic}} + \text{DM}_{\text{host}}, \quad (7)$$

where $\text{DM}_{\text{MW-disk}}$ is the contribution of the disk of the Milky Way, $\text{DM}_{\text{MW-halo}}$ is that from the extended hot Galactic halo and $\text{DM}_{\text{cosmic}}$ is from the intergalactic medium. The DM contribution of the host DM_{host} is a combination of the contributions from the interstellar medium of the host galaxy $\text{DM}_{\text{host-disk}}$, the halo of the host galaxy $\text{DM}_{\text{host-halo}}$ and the source environment $\text{DM}_{\text{host-env}}$.

To interpret unknown contributions to the total DM, we subtracted known contributions from the total. To estimate the contribution from the Milky Way disk, we defaulted to the NE2001 model^{20,21}, obtaining $\text{DM}_{\text{MW-disk, NE2001}} = (40 \pm 8)$ pc cm⁻³. Note that the YMW16 model⁵⁵ yielded similar results. We estimated the contribution of the Galactic halo to be $\text{DM}_{\text{MW-halo}} = (30 \pm 20)$ pc cm⁻³ using the model described in ref. 56. We treated this estimate as conservative, as it could be as low as 6 pc cm⁻³ (ref. 57). It is also consistent with CHIME/FRB constraints on the halo DM⁵⁸. The contribution of the intergalactic medium was estimated to be $\text{DM}_{\text{cosmic}} = (172 \pm 90)$ pc cm⁻³ (ref. 59), where the range is due to cosmic variance in the Macquart relation out to $z \approx 0.18$ (ref. 60). This

leaves the contribution to the DM from the host galaxy halo, disk and the FRB local environment as $\text{DM}_{\text{host}} = (257 \pm 93) \text{ pc cm}^{-3}$.

The large value of DM_{host} is consistent with a long line of sight travelled through the host galaxy disk, resulting from the galaxy inclination angle. We estimated the DM contributions of the host galaxy disk and halo by scaling the Milky Way's properties using the stellar mass of the host galaxy ('Host galaxy analysis' in Methods). We assumed that the disk size (R) scales with the galaxy stellar mass M_{host}^* as a power law $R \propto (M_{\text{host}}^*)^\beta$, where for simplicity we chose $\beta \approx 1/3$. This value of β is close to the measured value in the literature for galaxies with $M^* = 10^7 - 10^{11} M_\odot$ (ref. 61). Thus, the galaxy size scales as $(M_{\text{host}}^*/M_{\text{MW}}^*)^{1/3} = (1.4 \pm 0.3)^{1/3} = 1.12 \pm 0.08$, where $M_{\text{host}}^* = (8.5 \pm 0.8) \times 10^{10} M_\odot$ and $M_{\text{MW}}^* = (6.1 \pm 1.1) \times 10^{10} M_\odot$ are the present-day stellar masses of the Milky Way⁶² and the host galaxy, respectively. Assuming that the halo size also scales as $(M_{\text{host}}^*/M_{\text{MW}}^*)^{1/3}$, the average DM contribution of the Milky Way halo $(43 \pm 20) \text{ pc cm}^{-3}$ (ref. 56) can be scaled to estimate $\text{DM}_{\text{host-halo}}^r = \text{DM}_{\text{MW-halo}} \times (M_{\text{host}}^*/M_{\text{MW}}^*)^{1/3} = (48 \pm 23) \text{ pc cm}^{-3}$ in the host galaxy's rest frame. Similarly, we conservatively estimated the rest frame DM due to the disk of the host galaxy, $\text{DM}_{\text{host-disk}}^r$. A first approximation was to assume that the FRB originated from close to the midplane of the disk and to scale the DM contribution of the half-thickness of the Milky Way ($N_\perp(\infty) \approx (24 \pm 3) \text{ pc cm}^{-3}$)⁶³ by a factor of $\csc((7 \pm 3)^\circ) = 8 \pm 3$ to account for the viewing geometry. We assumed that the electron density was equivalent to that in the Milky Way and scaled for the host galaxy size. This yielded an estimate $\text{DM}_{\text{host-disk}}^r = N_\perp(\infty) \times \csc((7 \pm 3)^\circ) \times (M_{\text{host}}^*/M_{\text{MW}}^*)^{1/3} = (193 \pm 82) \text{ pc cm}^{-3}$ in the host galaxy rest frame. We summed these estimates of $\text{DM}_{\text{host-disk}}^r$ and $\text{DM}_{\text{host-halo}}^r$ to give the DM in the observer's frame as $\text{DM}_{\text{host}} = (\text{DM}_{\text{host-disk}}^r + \text{DM}_{\text{host-halo}}^r)/(1+z) = (224 \pm 82) \text{ pc cm}^{-3}$, which is consistent with the observed DM_{host} . If the FRB were behind the galaxy, the expected contribution from the host galactic disk could be increased by up to a factor of 2, yielding $(448 \pm 164) \text{ pc cm}^{-3}$; however, this possibility is inconsistent with the observed DM excess.

In addition to the DM of the host galaxy, we also measured gas fluctuations in the host galaxy using pulse broadening. The measured pulse broadening timescale from fitburst was $\tau_{\text{scatt-600 MHz}} = (165 \pm 3) \mu\text{s}$. However, after a visual inspection of the dynamic spectrum, we could not rule out the possibility that this timescale originated from an unresolved downward-drifting substructure. We treated this timescale as an upper limit on the true scattering timescale and considered the implications for FRB progenitors and the host galactic gas by comparing the dispersion and scattering to Galactic pulsars at similar Galactic latitudes. To place these measurements on an equal footing, we scaled $\tau_{\text{scatt-600 MHz}}$ to 1 GHz, and multiplied by $(1+z)^3$ to account for time dilation and the unredshifted frequency at which the pulse was scattered. This gave $\tau_{\text{proper,1GHz}} = 45 \mu\text{s}$ in the rest frame of the host galaxy. Dividing this by 3 converted the geometric weighting from extragalactic (plane-wave) scattering to Galactic (spherical-wave) scattering⁶⁴. Finally, subtracting $\text{DM}_{\text{host-halo}}$ from the observed DM excess in the host galaxy rest frame yielded $\text{DM}_{\text{host-disk}}^r = (254 \pm 111) \text{ pc cm}^{-3}$. We then calculated the ratio of observables:

$$\frac{\tau_{\text{proper,1GHz}}}{3(\text{DM}_{\text{host-disk}}^r)^2} \lesssim (4 \pm 3) \times 10^{-7} \text{ ms pc}^{-2} \text{ cm}^6 \propto FG.$$

This ratio characterizes the efficiency of the scattering along the line of sight. It is proportional to the product of the fluctuation parameter \bar{F} and an order-unity geometric factor G . The proportionality constant is $\Gamma(7/6)r_e^2 c^3 v^{-4}$, where $\Gamma(7/6) \approx 0.9277$, c is the speed of light, $r_e = 2.8 \text{ fm}$ is the classical electron radius and v is the frequency at which the scattering is observed⁶⁵. This proportionality constant captures the microphysics and the frequency dependence of the scattering and relates it to the ratio of observables. The bulk properties of the gas are captured by \bar{F} , which depends on the volume filling factor of gas

cloudlets, the size distribution of cloudlets doing the scattering, the size of the density variations within a cloudlet and the inner and outer scales of the turbulence⁶⁴. For the Milky Way's disk, typical values of \bar{F} range from 0.001 to $1 \text{ pc}^{-2/3} \text{ km}^{-1/3}$ for low-latitude sight lines, roughly corresponding to scattering DM^2 ratios of 10^{-8} to $10^{-5} \text{ ms pc}^{-2} \text{ cm}^6$ (ref. 64). G can vary by an order of magnitude because it depends on the relative position of the scattering media to the source and observer, which is poorly constrained for extragalactic sources of scattering. For example, for the geometry of a homogeneous scattering medium between the FRB and the edge of the host galaxy and a distant observer at infinity, $G = 1$. However, for a spiral arm of thickness $L \approx 1 \text{ kpc}$ at a distance $d \approx 10 \text{ kpc}$ in front of the FRB, $G = L/d \approx 0.1$. In conclusion, the host DM and scattering upper limit are consistent with expectations for a host galactic disk with Milky Way-like density fluctuations. These properties are suggestive of a source close to the host galaxy's plane as opposed to an FRB progenitor measurably displaced from the host galaxy's disk.

Another interpretation is that the DM excess is partially contributed by the source's local environment. The DM excess observed is not extreme. It is only a factor of two greater than the median measured in population studies ($\text{DM}_{\text{host}} \approx 145 \text{ pc cm}^{-3}$)⁶⁶. Furthermore, the upper limit on the scattering timescale and low RM are not outliers within the diverse population of FRBs. In this scenario, the FRB could have been produced by a progenitor displaced from the host galactic plane relative to the electron scale height $(1.57 \pm 0.15) \text{ kpc}$, which would reduce the host disk contribution to a fraction of our estimate $((224 \pm 82) \text{ pc cm}^{-3})$. This displacement would imply an old progenitor, as young progenitors typically have low scale heights, ~ 30 and 100 pc for young magnetars and massive stars, respectively^{67,68}.

Polarization analysis

The polarization analysis followed a similar procedure to that previously applied to other FRBs detected by CHIME/FRB^{69,70}. In particular, an initial RM estimate was made by applying RM-synthesis^{71,72} to the Stokes Q and U data of the burst. This initial estimate was then further refined through a judicious selection of time and frequency limits that optimized the S/N of the polarized signal. We then applied a Stokes QU -fitting routine that directly fitted for the modulation between Stokes Q and U from Faraday rotation but which had been extended to capture other features in the Stokes spectrum.

We analysed FRB 20210603A using the CHIME/FRB polarization pipeline, which is identical to that recently used for FRB 20191219F (ref. 73). We determined $\text{RM} = (-219.00 \pm 0.01) \text{ rad m}^{-2}$ and found that the lower limit of the linear polarized fraction (Π_l) differed between the top ($\geq 96\%$ at 800 MHz) and the bottom of the CHIME band ($\geq 87\%$ at 400 MHz). This was counteracted by a very small but changing circular polarized fraction that became more increased at the bottom of the band. Although this result may reflect the intrinsic properties of the burst at the source or be an imprint of some unknown propagation effect⁷⁴⁻⁷⁶, it is also not possible to rule out instrumental effects such as cross-polarization between CHIME's orthogonal feeds. For this reason, we do not report the circular polarization and conservatively set our Π_l measurements as lower bounds (Table 1).

The Galactic contribution $\text{RM}_{\text{MW}} = (-22.4 \pm 0.3) \text{ rad m}^{-2}$ was estimated from recent all-sky Faraday sky maps⁷⁷. The RM contribution of Earth's ionosphere, $\text{RM}_{\text{iono}} = +1.4 \text{ rad m}^{-2}$, was determined with the RMextract package⁷⁸. The uncertainty on this value was not provided. However, the variability in RM_{iono} was expected to be approximately less than $+1 \text{ rad m}^{-2}$ based on observations of pulsars and repeating FRB sources.

Given that the Galactic pulsar population preferentially occupies the Milky Way disk, this similarity, although it does not rule out alternative scenarios, is consistent with the notion that FRB 20210603A resides in or near the disk component of its host galaxy. Extended Data Fig. 10 explores this analysis by locating our DM_{host} , $|\text{RM}_{\text{host}}|$ and τ_{scatt} estimates

of FRB 20210603A within the equivalent phase space of the Galactic pulsar sample. Galactic pulsar data were obtained from the latest pulsar catalogue published by the Australia Telescope National Facility⁷⁹ using the `psrqpy` package⁸⁰. FRB 20210603A occupies a well-sampled region of this phase space. However, the distribution is highly dependent on the Galactic latitude. We estimated a quasi-latitude value for FRB 20210603A, determined from a simple transformation of the inclination angle of the host galaxy ($4^\circ \leq 90^\circ - \text{inclination angle} \leq 10^\circ$), and found that the average pulsar properties of DM, $|RM|$ and τ_{scatt} at this equivalent latitude agree well with what was observed for FRB 20210603A. The agreement was improved by rescaling DM and $|RM|$ to account for the larger disk mass of the host galaxy relative to the Milky Way. This scaling factor corresponds to the ratio of the disk masses of the host galaxy and the Milky Way, $(M_{\text{host}}^*/M_{\text{MW}}^*)^{1/3} = 1.12 \pm 0.08$ ('Dispersion and scattering analysis'). This result suggests that most of the DM_{host}, $|RM_{\text{host}}|$ and τ_{scatt} observed for FRB 20210603A could be supplied by the host galaxy's interstellar medium with little additional contribution needed from the source's local environment.

Host galaxy analysis

Optical images of SDSS J004105.82+211331.9 were taken with the CFHT MegaCam using the wideband *gri* filter. The data were reduced using the standard bias, dark and flat corrections in the Elixir pipeline^{81,82}. Several exposures were combined using this filter to create an image with a total exposure of 2,500 s.

The half-light radius of the host galaxy was determined using the Petrosian radii fluxes provided by SDSS Data Release 12 (ref. 8) and equation (7) of ref. 83. Using these values, the half-light radius in the *r*-band was (8.2 ± 0.9) kpc. Furthermore, the SDSS-provided apparent magnitude in the *r*-band was corrected for Milky Way extinction using the model from Fitzpatrick and Massa⁸⁴. This gave us an absolute magnitude of -22.03 ± 0.02 after *k*-corrections⁸⁵.

In addition to imaging, we conducted Gemini spectroscopic observations consisting of two 1,000 s exposures, one centred at 6,750 Å and the other at 6,650 Å. This wavelength offset was to account for the gap between the detectors. The images were reduced using standard bias and flat corrections and combined using the Gemini IRAF/PyRAF package tools^{86,87}. Using the same package, we also wavelength- and flux-calibrated the spectrum and accounted for skylines and cosmic rays in the data. We extracted spectra with various aperture sizes along the galaxy. The redshift was determined by extracting a spectrum from a 1 arcsec wide aperture centred at the central coordinates of the host galaxy. Due to the edge-on orientation of the galaxy, almost all of the galaxy's light falls within the slit, and the effect of slit corrections on the measured fluxes were negligible (Extended Data Fig. 8).

H α and the redwards line of the N II doublet (rest wavelengths of 6,564.6 and 6,585.2 Å) were some of the most detectable lines (Extended Data Fig. 8). Other prominent lines are from Na and Mg absorption (rest wavelengths of 5,895.6 and 5,176.7 Å). Fitting a linear combination of Gaussian line profiles to the H α and N II lines yielded a redshift of $z = 0.1772 \pm 0.0001$. The uncertainty in the spectroscopic redshift was dominated by the statistical uncertainties in the measured spectrum, which were normalized such that the reduced χ^2 of the residuals was 1.

To further characterize the galaxy, we combined our Gemini spectra with archival 2MASS (ref. 14) and WISE photometry¹⁵. We used the spectral-energy distribution fitting code Prospector to determine the stellar mass, metallicity and star formation history of the galaxy¹⁶. Our modelling and analysis of this host galaxy closely followed a similar effort for FRB 20181030A (ref. 88). However, because the galaxy is nearly edge on, dust extinction in the host galactic centre reddens the observed emission. Therefore, we first corrected the spectrum for extinction (equations (10) and (13) of ref. 89) due to its inclination of $(83 \pm 3)^\circ$ (ref. 10).

Our best-fitting model is overlaid on the spectral and photometric data in Extended Data Fig. 9. The model assumed a delayed- τ

star formation history $\propto t \exp(-t/\tau)$, where τ is the characteristic decay time and t is the time since the formation epoch of the galaxy. We set five free parameters: present-day stellar mass, metallicity, τ , t and the diffuse dust V-band optical depth (referred to as 'dust2' in Prospector), which accounted for the attenuation of old stellar light. We used τ and t as determined by Prospector to calculate the mass-weighted age of the galaxy. Additionally, we used a standard dust attenuation model⁹⁰ and enabled nebular emission and dust emission^{91,92}.

Before sampling the likelihood, we chose reasonable priors for each free parameter (Table 2). We used equation (6) of ref. 93 to obtain an initial estimate of the galaxy's mass and to set a weak prior on the mass range:

$$\log_{10}(M_{\text{host}}^*/M_{\odot}) = 1.097(g-r) - 4.06 - 0.4(M_r - 4.97) - 0.19z, \quad (8)$$

where g and r are the apparent magnitudes in the *g*-band and *r*-band filters, M_r is the absolute magnitude in the *r*-band and z is the redshift. The prior on t was cut off at 12 Gyr because the age of the Universe at $z = 0.1772 \pm 0.0001$ is only ~12 Gyr. The priors on Z/Z_{\odot} and τ were set according to recommendations in Prospector¹⁶. Using these priors, we obtained the fit plotted in Extended Data Fig. 9 and list the results in Table 1.

Finally, to determine the galaxy-integrated SFR, we extracted a spectrum with an aperture of 10 arcsec in diameter, which encompassed all of the galaxy's light within our half-light radius of ~2.5 arcsec. We calculated the total SFR of the host galaxy using the intensity and linewidth of the H α line⁹⁴:

$$\text{SFR} = 7.9e - 42 \left(\frac{L_{\text{H}\alpha}}{\text{erg s}^{-1}} \right) \frac{M_{\odot}}{\text{yr}}, \quad (9)$$

where $L_{\text{H}\alpha}$ is the flux-derived luminosity of the H α emission from our Gemini data. To correct our luminosity measurement for extinction, we applied the inclination-dependent correction as well as the inclination-independent correction, parameterized as dust2 in Prospector. The latter quantifies the amount of V-band extinction of old stellar light in the host galaxy. Optical reddening was characterized using $R_v = A_v/E(B-V)$, where $E(B-V)$ is the colour index of the galaxy and A_v is the extinction in the V-band. This equation is, thus, the ratio of total to selective extinction in the V-band⁹⁵. The dust extinction was taken to be $A_v = 1.086 \times \text{dust2}$ (ref. 96), where we took dust2 to be the best-fitting value of 0.79. With $R_v = 3.1$ (ref. 95), we calculated $E(B-V)$ to be 0.28. The H α extinction coefficient was calculated using $A_{\text{H}\alpha} = R_{\text{H}\alpha} \times E(B-V)$, where we took $R_{\text{H}\alpha} = 2.45$, which is within the range of values predicted by several different extinction models^{84,90,97,98}. The inclination-independent attenuation resulted in the H α flux being attenuated by a factor of $\exp(A_{\text{H}\alpha}) = 1.97$. Correcting the galaxy-integrated H α flux for extinction yielded a total SFR of $0.24 \pm 0.06 M_{\odot} \text{ yr}^{-1}$.

Disk chance coincidence probability

Although FRB 20210603A was ostensibly localized to the disk of its host galaxy, it is possible that the progenitor is actually a halo object (like the globular cluster host of FRB 20200120E; ref. 99) coincidentally aligned with the disk in projection. The probability that this occurs by a chance coincidence is small. We estimated this probability as the ratio of the solid angles subtended by the disk and halo, $P_{\text{cc}} \approx \Omega_{\text{disk}}/\Omega_{\text{halo}} \approx 10^{-3}$. The angular area of the nearly edge-on disk was approximated as an ellipse with major and minor axes of 15 and 2.7 arcsec, respectively, whereas the area of the halo was approximated as a circle of radius $r_{\text{vir}} \approx M_{\text{host}}^*/M_{\text{MW}}^* r_{\text{vir,MW}} \approx 280$ kpc, which was estimated by scaling up the Milky Way's virial radius $r_{\text{vir,MW}} \approx 200$ kpc (ref. 100). This low chance coincidence probability of 10^{-3} implies a robust association with the disk and favours progenitor models involving disk populations over halo populations.

Data availability

Calibrated visibilities, dynamic spectra for producing the figures and Markov-chain Monte Carlo chains for the localization analysis are available upon request and will be hosted by the time of publication as downloadable HDF5 files from the repository of the Canadian Advanced Network for Astronomical Research for CHIME/FRB at <https://www.canfar.net/storage/list/AstroDataCitationDOI/CISTI.CANFAR/24.0086/data>. Optical images, spectra and photometric data are immediately available as fits files at <https://github.com/tcassanelli/frb-vlbi-loc>.

Code availability

The code used for beamforming, VLBI localization and polarization analysis are available on Github: <https://github.com/CHIMEFRB/baseband-analysis>. The scattering timescale was measured using `fitburst`⁵⁴, which is available at <https://github.com/CHIMEFRB/fitburst>. Code for interpreting burst properties and for producing the figures and tables in this paper from the results of our analyses is available at <https://github.com/tcassanelli/frb-vlbi-loc>. In our analyses, we also made use of open-source software including `Astropy`¹⁰¹, `baseband`¹⁰², `difxcalc11`⁴³, `Matplotlib`¹⁰³, `NumPy`¹⁰⁴, `SciPy`¹⁰⁵, `h5py`¹⁰⁶, `emcee`¹⁷, `corner`¹⁰⁷, `cartopy`¹⁰⁸, `IRAF`^{86,87} and `Prospector`¹⁶.

References

- CHIME/FRB Collaboration et al. The CHIME fast radio burst project: system overview. *Astrophys. J.* **863**, 48 (2018).
- CHIME Collaboration et al. An overview of CHIME, the Canadian Hydrogen Intensity Mapping Experiment. *Astrophys. J. Suppl. Ser.* **261**, 29 (2022).
- Cassanelli, T. et al. Localizing FRBs through VLBI with the Algonquin Radio Observatory 10 m telescope. *Astron. J.* **163**, 65 (2022).
- Sanghavi, P. et al. TONE: a CHIME/FRB outtrigger pathfinder for localizations of fast radio bursts using very long baseline interferometry. Preprint at arxiv.org/abs/2304.10534 (2023).
- Michilli, D. et al. An analysis pipeline for CHIME/FRB full-array baseband data. *Astrophys. J.* **910**, 147 (2021).
- Leung, C. et al. A VLBI software correlator for fast radio transients. Preprint at arxiv.org/abs/2403.05631 (2024).
- Eubanks, M. et al. In *Proc. US Naval Observatory Workshop on Relativistic Models for use in Space Geodesy* 60–82 (US Naval Observatory, 1991).
- Ahn, C. P. et al. The tenth data release of the Sloan Digital Sky Survey: first spectroscopic data from the SDSS-III Apache Point Observatory Galactic Evolution Experiment. *Astrophys. J. Suppl. Ser.* **211**, 17 (2014).
- Boulade, O. et al. MegaCam: the next-generation wide-field imaging camera for CFHT. In *Proc. SPIE Conference Series, Optical Astronomical Instrumentation* Vol. 3355 (ed. D'Odorico, S.) 614–625 (SPIE, 1998).
- Kourkchi, E. et al. Cosmicflows-4: the catalog of ~10,000 Tully–Fisher distances. *Astrophys. J.* **902**, 145 (2020).
- Hook, I. M. et al. The Gemini-north multi-object spectrograph: performance in imaging, long-slit, and multi-object spectroscopic modes. *Publ. Astron. Soc. Pac.* **116**, 425–440 (2004).
- Planck Collaboration et al. Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.* **641**, A6 (2020).
- Bhandari, S. et al. Characterizing the fast radio burst host galaxy population and its connection to transients in the local and extragalactic Universe. *Astron. J.* **163**, 69 (2022).
- Skrutskie, M. F. et al. The Two Micron All Sky Survey (2MASS). *Astron. J.* **131**, 1163–1183 (2006).
- Wright, E. L. et al. The Wide-field Infrared Survey Explorer (WISE): mission description and initial on-orbit performance. *Astron. J.* **140**, 1868–1881 (2010).
- Johnson, B. D., Leja, J., Conroy, C. & Speagle, J. S. Stellar population inference with Prospector. *Astrophys. J. Suppl. Ser.* **254**, 22 (2021).
- Foreman-Mackey, D., Hogg, D. W., Lang, D. & Goodman, J. emcee: the MCMC hammer. *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- Niu, C. H. et al. A repeating fast radio burst associated with a persistent radio source. *Nature* **606**, 873–877 (2022).
- CHIME/FRB Collaboration et al. The First CHIME/FRB fast radio burst catalog. *Astrophys. J. Suppl. Ser.* **257**, 59 (2021).
- Cordes, J. M. & Lazio, T. J. W. NE2001.I. A new model for the Galactic distribution of free electrons and its fluctuations. Preprint at arxiv.org/abs/astro-ph/0207156 (2002).
- Cordes, J. M. & Lazio, T. J. W. NE2001.II. Using radio propagation data to construct a model for the Galactic distribution of free electrons. Preprint at arxiv.org/abs/astro-ph/0301598 (2003).
- Masui, K. et al. Dense magnetized plasma associated with a fast radio burst. *Nature* **528**, 523–525 (2015).
- Akahori, T., Ryu, D. & Gaensler, B. M. Fast radio bursts as probes of magnetic fields in the intergalactic medium. *Astrophys. J.* **824**, 105 (2016).
- Tendulkar, S. P. et al. The 60 pc environment of FRB 20180916B. *Astrophys. J. Lett.* **908**, L12 (2021).
- Leung, C. et al. A synoptic VLBI technique for localizing nonrepeating fast radio bursts with CHIME/FRB. *Astron. J.* **161**, 81 (2021).
- Mena-Parra, J. et al. A clock stabilization system for CHIME/FRB outtriggers. *Astron. J.* **163**, 48 (2022).
- Bandura, K. et al. ICE: a scalable, low-cost FPGA-based telescope signal processing and networking system. *J. Astron. Instrum.* **5**, 1641005 (2016).
- Denman, N. et al. A GPU spatial processing system for CHIME. *J. Astron. Instrum.* **9**, 2050014–2 (2020).
- Renard, A. et al. Kotekan: A framework for high-performance radiometric data pipelines. Zenodo <https://zenodo.org/records/5842660> (2021).
- Ng, C. et al. CHIME/FRB: an application of FFT beamforming for a radio telescope. In *Proc. 2017 XXXIInd General Assembly and Scientific Symposium of the International Union of Radio Science (URSI GASS)* 1–4 (IEEE, 2017).
- Lyne, A. G. et al. 45 years of rotation of the Crab pulsar. *Mon. Not. R. Astron. Soc.* **446**, 857–864 (2015).
- Deng, M. & Campbell-Wilson, D. The cloverleaf antenna: a compact wide-bandwidth dual-polarization feed for CHIME. In *Proc. 2014 16th International Symposium on Antenna Technology and Applied Electromagnetics (ANTEN)* 1–2 (IEEE, 2014).
- Crichton, D. et al. Hydrogen intensity and real-time analysis experiment: 256-element array status and overview. *J. Astron. Telesc. Instrum. Syst.* **8**, 011019 (2022).
- Mena, J. et al. A radio-frequency-over-fiber link for large-array radio astronomy applications. *J. Instrum.* **8**, T10003–T10003 (2013).
- Spectrum Instruments, I. *Intelligent Reference/TM-4* www.spectruminstruments.net/products/tm4/tm4.html (2008).
- Cary, S. et al. Evaluating and enhancing candidate clocking systems for CHIME/FRB VLBI outtriggers. *Res. Notes AAS* **5**, 216 (2021).
- Sanghavi, P. R. *Pathfinding Fast Radio Bursts Localizations Using Very Long Baseline Interferometry*. PhD thesis, West Virginia Univ. (2022).
- Cassanelli, T. *Fast Radio Burst Localization with Very Long Baseline Interferometry*. PhD thesis, Univ. of Toronto (2022).
- Leung, C. *Localization and Lensing of Fast Radio Bursts using CHIME/FRB and its VLBI Outtriggers*. PhD thesis, Massachusetts Institute of Technology (2023).

40. Recnik, A. et al. An efficient real-time data pipeline for the CHIME pathfinder radio telescope X-Engine. In *Proc. 2015 IEEE 26th International Conference on Application-specific Systems, Architectures and Processors (ASAP)* (eds Anderson, J. et al.) (IEEE, 2015).
41. Bandura, K. et al. Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder. In *Proc. SPIE Conference Series, Ground-based and Airborne Telescopes V* Vol. 9145 (eds Stepp, L. M. et al.) 22 (SPIE, 2014).
42. Newburgh, L. B. et al. Calibrating CHIME: a new radio interferometer to probe dark energy. In *Proc. SPIE Conference Series, Ground-based and Airborne Telescopes V* Vol. 9145 (eds Stepp, L. M. et al.) 4V (SPIE, 2014).
43. Gordon, D., Briske, W. & Max-Moerbeck, W. Difxcalc - Calc11 for the DiFX Correlator. In *Proc. New Horizons with VGOS* (eds Behrend, D. et al.) 187–192 (NASA, 2016).
44. Lobanov, A. P., Horns, D. & Muxlow, T. W. B. VLBI imaging of a flare in the Crab nebula: more than just a spot. *Astron. Astrophys.* **533**, A10 (2011).
45. Hankins, T. H. & Rickett, B. J. Pulsar signal processing. *Methods Comput. Phys.: Adv. Res. Appl.* **14**, 55–129 (1975).
46. Lorimer, D. R. & Kramer, M. *Handbook of Pulsar Astronomy* (Cambridge Univ. Press, 2012).
47. Kulkarni, S. R. Dispersion measure: confusion, constants & clarity. Preprint at arxiv.org/abs/2007.02886 (2020).
48. Nimmo, K. et al. Milliarsecond localization of the repeating FRB 20201124A. *Astrophys. J. Lett.* **927**, L3 (2022).
49. Ma, C. et al. The International Celestial Reference Frame as realized by very long baseline interferometry. *Astron. J.* **116**, 516–546 (1998).
50. Brown, A. G. A. Microarcsecond astrometry: science highlights from Gaia. *Annu. Rev. Astron. Astrophys.* **59**, 59–115 (2021).
51. Kaplan, G. H. *The IAU Resolutions on Astronomical Reference Systems, Time Scales, and Earth Rotation Models: Explanation and Implementation* Circular No. 179 (US Naval Observatory, 2005).
52. Feissel, M. & Mignard, F. The adoption of ICRS on 1 January 1998: meaning and consequences. *Astron. Astrophys.* **331**, L33–L36 (1998).
53. Seymour, A., Michilli, D. & Pleunis, Z. DM_phase: algorithm for correcting dispersion of radio signals. Astrophysics Source Code Library, record ascl:1910.004 1910.004 (ASCL, 2019).
54. Fonseca, E. et al. Modeling the morphology of fast radio bursts and radio pulsars with fitburst. *Astrophys. J. Suppl. Ser.* **271**, 49 (2024).
55. Yao, J. M., Manchester, R. N. & Wang, N. A new electron-density model for estimation of pulsar and FRB distances. *Astrophys. J.* **835**, 29 (2017).
56. Yamasaki, S. & Totani, T. The Galactic halo contribution to the dispersion measure of extragalactic fast radio bursts. *Astrophys. J.* **888**, 105 (2020).
57. Keating, L. C. & Pen, U.-L. Exploring the dispersion measure of the Milky Way halo. *Mon. Not. R. Astron. Soc.* **496**, L106–L110 (2020).
58. Cook, A. M. et al. An FRB sent me a DM: constraining the electron column of the Milky Way Halo with fast radio burst dispersion measures from CHIME/FRB. *Astrophys. J.* **946**, 58 (2023).
59. Macquart, J. P. et al. A census of baryons in the Universe from localized fast radio bursts. *Nature* **581**, 391–395 (2020).
60. Batten, A. J. et al. The cosmic dispersion measure in the EAGLE simulations. *Mon. Not. R. Astron. Soc.* **505**, 5356–5369 (2021).
61. Trujillo, I., Chamba, N. & Knapen, J. H. A physically motivated definition for the size of galaxies in an era of ultradeep imaging. *Mon. Not. R. Astron. Soc.* **493**, 87–105 (2020).
62. Licquia, T. C. & Newman, J. A. Improved estimates of the Milky Way's stellar mass and star formation rate from hierarchical Bayesian meta-analysis. *Astrophys. J.* **806**, 96 (2015).
63. Ocker, S. K., Cordes, J. M. & Chatterjee, S. Electron density structure of the local Galactic disk. *Astrophys. J.* **897**, 124 (2020).
64. Ocker, S. K., Cordes, J. M. & Chatterjee, S. Constraining galaxy halos from the dispersion and scattering of fast radio bursts and pulsars. *Astrophys. J.* **911**, 102 (2021).
65. Cordes, J. M., Wharton, R. S., Spitler, L. G., Chatterjee, S. & Wasserman, I. Radio wave propagation and the provenance of fast radio bursts. Preprint at arxiv.org/abs/1605.05890 (2016).
66. James, C. W. et al. The z-DM distribution of fast radio bursts. *Mon. Not. R. Astron. Soc.* **509**, 4775–4802 (2022).
67. Olausen, S. A. & Kaspi, V. M. The McGill magnetar catalog. *Astrophys. J. Suppl. Ser.* **212**, 6 (2014).
68. Miller, G. E. & Scalo, J. M. The initial mass function and stellar birthrate in the solar neighborhood. *Astrophys. J. Suppl. Ser.* **41**, 513 (1979).
69. Fonseca, E. et al. Nine new repeating fast radio burst sources from CHIME/FRB. *Astrophys. J. Lett.* **891**, L6 (2020).
70. Bhardwaj, M. et al. A nearby repeating fast radio burst in the direction of M81. *Astrophys. J. Lett.* **910**, L18 (2021).
71. Burn, B. J. On the depolarization of discrete radio sources by Faraday dispersion. *Mon. Not. R. Astron. Soc.* **133**, 67 (1966).
72. Brentjens, M. A. & de Bruyn, A. G. Faraday rotation measure synthesis. *Astron. Astrophys.* **441**, 1217–1228 (2005).
73. McKinven, R. et al. Polarization pipeline for fast radio bursts detected by CHIME/FRB. *Astrophys. J.* **920**, 138 (2021).
74. Vedantham, H. K. & Ravi, V. Faraday conversion and magneto-ionic variations in fast radio bursts. *Mon. Not. R. Astron. Soc.* **485**, L78–L82 (2019).
75. Gruzinov, A. & Levin, Y. Conversion measure of Faraday rotation-conversion with application to fast radio bursts. *Astrophys. J.* **876**, 74 (2019).
76. Beniamini, P., Kumar, P. & Narayan, R. Faraday depolarization and induced circular polarization by multipath propagation with application to FRBs. *Mon. Not. R. Astron. Soc.* **510**, 4654–4668 (2022).
77. Hutschenreuter, S. et al. The Galactic Faraday rotation sky 2020. *Astron. Astrophys.* **657**, A43 (2022).
78. Mevius, M. RMextract: ionospheric Faraday rotation calculator. Astrophysics Source Code Library, record ascl:1806.024 (ASCL, 2018).
79. Manchester, R. N., Hobbs, G. B., Teoh, A. & Hobbs, M. The Australia Telescope National Facility pulsar catalogue. *Astron. J.* **129**, 1993–2006 (2005).
80. Pitkin, M. psrqpy: a Python interface for querying the ATNF pulsar catalogue. *J. Open Source Softw.* **3**, 538 (2018).
81. Magnier, E. A. & Cuillandre, J. C. The Elixir System: data characterization and calibration at the Canada–France–Hawaii Telescope. *Publ. Astron. Soc. Pac.* **116**, 449–464 (2004).
82. Prunet, S., Fouque, P. & Gwyn, S. *Photometric Calibration of MegaCam Data* <https://www.cfht.hawaii.edu/Instruments/Imaging/MegaPrime/PDFs/megacam.pdf> (2014).
83. Graham, A. W. et al. Total galaxy magnitudes and effective radii from Petrosian magnitudes and radii. *Astron. J.* **130**, 1535–1544 (2005).
84. Fitzpatrick, E. L. & Massa, D. An analysis of the shapes of interstellar extinction curves. V. The IR-through-UV curve morphology. *Astrophys. J.* **663**, 320–341 (2007).
85. Chilingarian, I. V., Melchior, A.-L. & Zolotukhin, I. Y. Analytical approximations of K-corrections in optical and near-infrared bands. *Mon. Not. R. Astron. Soc.* **405**, 1409–1420 (2010).
86. Tody, D. The IRAF data reduction and analysis system. In *Proc. SPIE Conference Series, Instrumentation in Astronomy VI* Vol. 627 (ed. Crawford, D. L.) 733 (SPIE, 1986).
87. Tody, D. in *Astronomical Data Analysis Software and Systems II* Vol. 52 (eds Hanisch, R. J. et al.) 173 (ASP, 1993).

88. Bhardwaj, M. et al. A local Universe host for the repeating fast radio burst FRB 20181030a. *Astrophys. J. Lett.* **919**, L24 (2021).
89. Shao, Z. et al. Inclination-dependent luminosity function of spiral galaxies in the Sloan Digital Sky Survey: implications for dust extinction. *Astrophys. J.* **659**, 1159–1171 (2007).
90. Calzetti, D. et al. The dust content and opacity of actively star-forming galaxies. *Astrophys. J.* **533**, 682–695 (2000).
91. Byler, N., Dalcanton, J. J., Conroy, C. & Johnson, B. D. Nebular continuum and line emission in stellar population synthesis models. *Astrophys. J.* **840**, 44 (2017).
92. Draine, B. T. & Li, A. Infrared emission from interstellar dust. IV. The silicate-graphite-PAH model in the post-Spitzer era. *Astrophys. J.* **657**, 810–837 (2007).
93. Bernardi, M. et al. Galaxy luminosities, stellar masses, sizes, velocity dispersions as a function of morphological type. *Mon. Not. R. Astron. Soc.* **404**, 2087–2122 (2010).
94. Kennicutt Jr, R. C., Tamblyn, P. & Congdon, C. E. Past and future star formation in disk galaxies. *Astrophys. J.* **435**, 22 (1994).
95. Fitzpatrick, E. L. Correcting for the effects of interstellar extinction. *Publ. Astron. Soc. Pac.* **111**, 63–75 (1999).
96. Conroy, C., Gunn, J. E. & White, M. The propagation of uncertainties in stellar population synthesis modeling. I. The relevance of uncertain aspects of stellar evolution and the initial mass function to the derived physical properties of galaxies. *Astrophys. J.* **699**, 486–506 (2009).
97. Schlafly, E. F. & Finkbeiner, D. P. Measuring reddening with Sloan Digital Sky Survey stellar spectra and recalibrating SFD. *Astrophys. J.* **737**, 103 (2011).
98. O'Donnell, J. E. R_v -dependent optical and near-ultraviolet extinction. *Astrophys. J.* **422**, 158 (1994).
99. Kirsten, F. et al. A repeating fast radio burst source in a globular cluster. *Nature* **602**, 585–589 (2022).
100. Dehnen, W., McLaughlin, D. E. & Sachania, J. The velocity dispersion and mass profile of the Milky Way. *Mon. Not. R. Astron. Soc.* **369**, 1688–1692 (2006).
101. Astropy Collaboration et al. The Astropy project: building an open-science project and status of the v2.0 core package. *Astron. J.* **156**, 123 (2018).
102. van Kerkwijk, M. et al. mshv/baseband: v4.1.3. *Zenodo* <https://doi.org/10.5281/zenodo.7941170> (2023).
103. Hunter, J. D. Matplotlib: a 2D graphics environment. *Comput. Sci. Eng.* **9**, 90–95 (2007).
104. Harris, C. R. et al. Array programming with NumPy. *Nature* **585**, 357–362 (2020).
105. Virtanen, P. et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**, 261–272 (2020).
106. Collette, A. et al. h5py/h5py: 3.2.1. *Zenodo* <https://zenodo.org/record/4584676> (2021).
107. Foreman-Mackey, D. corner.py: scatterplot matrices in Python. *J. Open Source Softw.* **1**, 24 (2016).
108. Met Office. Cartopy: A Cartographic Python Library with a Matplotlib Interface. <http://scitools.org.uk/cartopy> (2015).

Acknowledgements

We would like to dedicate this work to our colleague Jing Luo, who passed away on 15 February 2022. Jing worked for several years on the commissioning and maintenance of the 10m telescope at Algonquin Radio Observatory. His expertise in radio pulsars, radio observations and hardware were indispensable in accomplishing this scientific milestone. We acknowledge that CHIME is located on the traditional, ancestral and unceded territory of the Syilx/Okanagan people. CHIME is funded by a grant from the 2012 Leading Edge Fund of the Canada Foundation for Innovation (Project No. 31170) and by contributions from the provinces of British Columbia, Québec and Ontario.

The CHIME/FRB Project is funded by a grant from the 2015 Innovation Fund of the Canada Foundation for Innovation (Project No. 33213) and by contributions from the provinces of British Columbia and Québec, and by the Dunlap Institute for Astronomy and Astrophysics at the University of Toronto. Further support was provided by the Canadian Institute for Advanced Research (CIFAR), McGill University and the McGill Space Institute thanks to the Trottier Family Foundation and the University of British Columbia. The Dunlap Institute is funded through an endowment established by the David Dunlap family and the University of Toronto. We thank the staff of DRAO, which is operated by the National Research Council of Canada, for their gracious hospitality, support and expertise. ARO10 is operated by the University of Toronto. TONE is at Green Bank Observatory, which is supported by the National Science Foundation (NSF), and is operated by Associated Universities, Inc. under a cooperative agreement. We would like to thank the staff at Green Bank Observatory for logistical support during the construction and operations of TONE. This work is based on observations obtained at the international Gemini Observatory, a programme of NSF's NOIRLab, which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the NSF. PyRAF is a product of the Space Telescope Science Institute, which is operated by AURA for NASA on behalf of the Gemini Observatory partnership: the NSF (United States), National Research Council (Canada), Agencia Nacional de Investigación y Desarrollo (Chile), Ministerio de Ciencia, Tecnología e Innovación (Argentina), Ministério da Ciência, Tecnologia, Inovações e Comunicações (Brazil) and the Korea Astronomy and Space Science Institute (Republic of Korea). This work is also based on observations obtained with MegaPrime/MegaCam, a joint project of CFHT and the French Alternative Energies and Atomic Energy Commission/Dapnia, at the CFHT, which is operated by the National Research Council of Canada, the Institut National des Sciences de l'Univers of the French National Centre for Scientific Research and the University of Hawaii. The observations at the CFHT were performed with care and respect from the summit of Maunakea which is a significant cultural and historic site. A.B.P. is a Banting Fellow, a McGill Space Institute Fellow and a Fonds de Recherche du Québec – Nature et Technologies (FRQNT) postdoctoral fellow. A.P.C. is a Vanier Canada Graduate Scholar. B.M.G. acknowledges the support of the Natural Sciences and Engineering Research Council of Canada (NSERC; Grant No. RGPIN-2015-05948) and of the Canada Research Chairs programme. C.L. was supported by the US Department of Defense through the National Defense Science & Engineering graduate fellowship programme and is an NHFP Einstein Fellow. E.P. acknowledges funding from a Veni Fellowship from the Netherlands Organisation for Scientific Research. F.A.D. is funded by a four-year doctoral fellowship at the University of British Columbia. FRB research at the University of British Columbia is funded by an NSERC discovery grant and by the Canadian Institute for Advanced Research. The CHIME baseband system was funded in part by an award granted by the John R. Evans Leaders Fund from the Canada Foundation for Innovation to I.S. FRB research at West Virginia University is supported by the NSF (Grant Nos. 2006548 and 2018490). J.B.P. is supported by the Major Research Instrumentation Program of the NSF (Grant No. 2018490). J.M.P. is a Kavli Fellow. K. Shin is supported by the NSF Graduate Research Fellowship Program. K.W.M. is supported by the NSF (Grant No. 2008031). M.B. is supported by an FRQNT doctoral research award. M.D. is supported by a Canada Research Chair, Killam Fellowship and NSERC Discovery Grant, by CIFAR and by the FRQNT Centre de Recherche en Astrophysique du Québec. P. Scholz is a Dunlap Fellow. S. Cary would like to thank K. McLeod from Wellesley College for her supervision and feedback, which was essential for the host galaxy analysis. S. Chatterjee is a member of the NANOGrav Physics Frontiers Center, which is supported by the NSF (Award No. PHY-1430284). U.P. is supported by NSERC (Funding References RGPIN-2019-067,

CRD 523638-201 and 555585-20), the Research Excellence Program of the Ontario Research Fund, CIFAR, the Simons Foundation, Thoth Technology Inc. and the Alexander von Humboldt Foundation. This research is supported by the Ministry of Science and Technology of Taiwan (Grant No. 110-2112-M-001-071-MY3). V.M.K. holds the Lorne Trottier Chair in Astrophysics & Cosmology and a Distinguished James McGill Professorship and receives support from an NSERC Discovery grant (RGPIN 228738-13), an R. Howard Webster Foundation Fellowship from CIFAR and the FRQNT Centre de Recherche en Astrophysique du Québec. Z.P. is a Dunlap Fellow.

Author contributions

T.C. led the full instrument adaptation, triggering system and data management of the ARO10 telescope, wrote the main text and methods sections, and prepared several figures and tables. C.L. wrote the VLBI software correlator and analysis pipeline used to localize the FRB, designed and built the digital back end of the TONE array, and led the data analysis, scientific interpretation and writing of the paper. P. Sanghavi. led the design, construction, commissioning and data acquisition of all aspects of the TONE telescope and contributed significantly to the scientific interpretation and writing of the paper. J.M.P. designed and installed the maser hardware, characterized the clock stabilization system and made foundational contributions to the array calibration pipelines used at TONE and CHIME. S. Cary led the calibration, reduction and analysis of the optical follow-up data. All other authors from the CHIME/FRB collaboration played either leadership or significant supporting roles in one or more of the following activities: management, development, construction, commissioning and maintenance of CHIME, the CHIME/FRB instrument, the ARO10 instrument, the TONE instrument,

their respective software data pipelines, or the data analysis and preparation of this paper.

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at <https://doi.org/10.1038/s41550-024-02357-x>.

Correspondence and requests for materials should be addressed to Calvin Leung.

Peer review information *Nature Astronomy* thanks Adam Deller and Zsolt Paragi for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

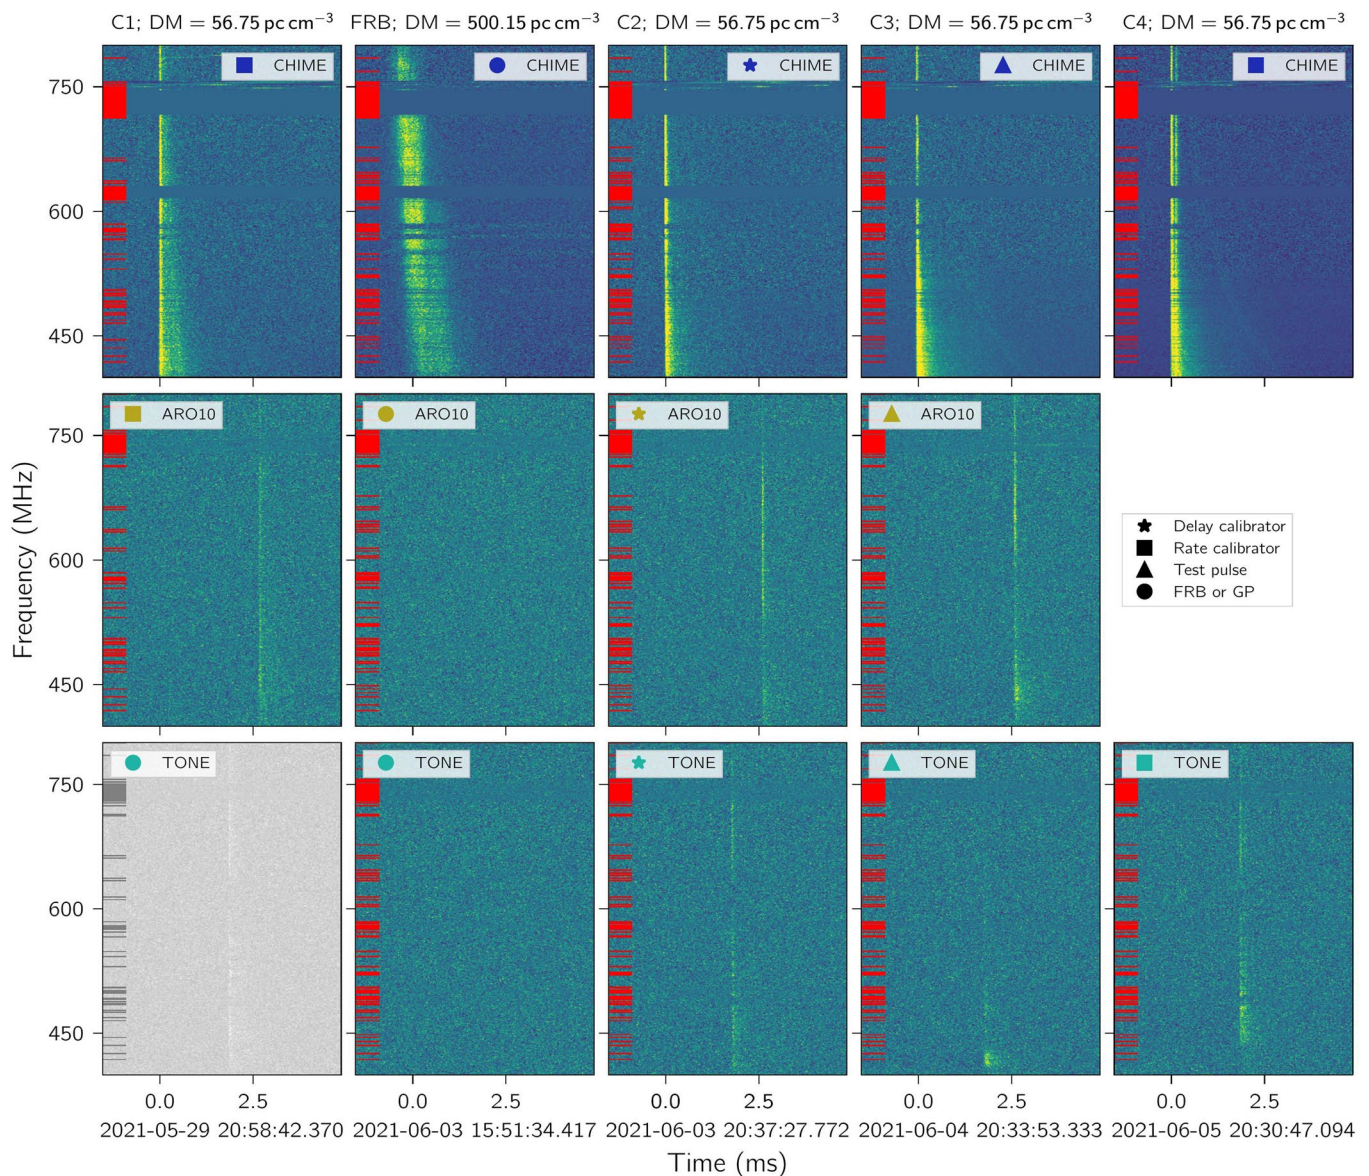
Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© The Author(s), under exclusive licence to Springer Nature Limited 2024


Tomas Cassanelli ^{1,2,3,31}, **Calvin Leung** ^{4,5,31} ✉, **Pranav Sanghavi** ^{6,7,8,31}, **Juan Mena-Parra** ^{1,2}, **Savannah Cary** ^{4,9}, **Ryan Mckinven** ^{10,11}, **Mohit Bhardwaj** ^{10,11}, **Kiyoshi W. Masui** ^{4,5}, **Daniele Michilli** ^{4,5}, **Kevin Bandura** ^{6,7}, **Shami Chatterjee** ¹², **Jeffrey B. Peterson** ¹³, **Jane Kaczmarek** ^{14,15,16}, **Mubdi Rahman** ¹⁷, **Kaitlyn Shin** ^{4,5}, **Keith Vanderlinde** ^{1,2}, **Sabrina Berger** ^{10,11}, **Charanjot Brar** ^{10,11}, **P. J. Boyle** ^{10,11}, **Daniela Breitman** ¹⁸, **Pragya Chawla** ¹⁹, **Alice P. Curtin** ^{10,11}, **Matt Dobbs** ^{10,11}, **Fengqiu Adam Dong** ²⁰, **Emmanuel Fonseca** ^{7,21}, **B. M. Gaensler** ^{1,2,22}, **Adaeze Ibik** ^{1,2}, **Victoria M. Kaspi** ^{10,11}, **Kholoud Khairy** ^{6,7}, **Adam E. Lanman** ^{10,11}, **Mattias Lazda** ^{1,10}, **Hsiu-Hsien Lin** ^{23,24}, **Jing Luo** ^{24,32,33}, **Bradley W. Meyers** ²⁰, **Nikola Milutinovic** ²⁰, **Cherry Ng** ², **Gavin Noble** ^{1,2}, **Aaron B. Pearlman** ^{10,11}, **Ue-Li Pen** ^{2,4,24,25,26}, **Emily Petroff** ^{10,11,19}, **Ziggy Pleunis** ², **Brendan Quine** ^{27,28}, **Masoud Rafiei-Ravandi** ^{10,11}, **Andre Renard** ², **Ketan R. Sand** ^{10,11}, **Eve Schoen** ⁴, **Paul Scholz** ², **Kendrick M. Smith** ²⁶, **Ingrid Stairs** ²⁰ & **Shriharsh P. Tendulkar** ^{29,30}

¹David A. Dunlap Department of Astronomy & Astrophysics, University of Toronto, Toronto, Ontario, Canada. ²Dunlap Institute for Astronomy & Astrophysics, University of Toronto, Toronto, Ontario, Canada. ³Department of Electrical Engineering, Universidad de Chile, Santiago, Chile. ⁴MIT Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA, USA. ⁵Department of Physics, Massachusetts Institute of Technology, Cambridge, MA, USA. ⁶Lane Department of Computer Science and Electrical Engineering, West Virginia University, Morgantown, WV, USA. ⁷Center for Gravitational Waves and Cosmology, West Virginia University, Morgantown, WV, USA. ⁸Department of Physics, Yale University, New Haven, CT, USA. ⁹Department of Astronomy, Wellesley College, Wellesley, MA, USA. ¹⁰Department of Physics, McGill University, Montreal, Quebec, Canada. ¹¹Trottier Space Institute, McGill University, Montreal, Quebec, Canada. ¹²Cornell Center for Astrophysics and Planetary Science, Ithaca, NY, USA. ¹³McWilliams Center for Cosmology, Department of Physics, Carnegie Mellon University, Pittsburgh, PA, USA. ¹⁴Dominion Radio Astrophysical Observatory, Herzberg Research Centre for Astronomy and Astrophysics, National Research Council Canada, Penticton, British Columbia, Canada. ¹⁵Department of Computer Science, Math, Physics, & Statistics, University of British Columbia, Kelowna, British Columbia, Canada. ¹⁶CSIRO Space & Astronomy, Parkes Observatory, Parkes, New South Wales, Australia. ¹⁷Sidrat Research, Toronto, Ontario, Canada. ¹⁸Scuola Normale Superiore, Pisa, Italy. ¹⁹Anton Pannekoek Institute for Astronomy, University of Amsterdam, Amsterdam, The Netherlands. ²⁰Department of Physics and Astronomy, University of British Columbia, Vancouver, British Columbia, Canada. ²¹Department of Physics and Astronomy, West Virginia University, Morgantown, WV, USA. ²²Department of Astronomy and Astrophysics, University of California Santa Cruz, Santa Cruz, CA, USA. ²³Institute of Astronomy and Astrophysics, Academia Sinica, Taipei, Taiwan. ²⁴Canadian Institute for Theoretical Astrophysics, Toronto, Ontario, Canada. ²⁵Canadian Institute for Advanced Research, Toronto, Ontario, Canada. ²⁶Perimeter Institute for Theoretical Physics, Waterloo, Ontario, Canada. ²⁷Thoth Technology Inc., Deep River, Ontario, Canada. ²⁸Department of Physics and Astronomy, York University, Toronto, Ontario, Canada. ²⁹Department of Astronomy and Astrophysics, Tata Institute of Fundamental Research, Mumbai, India. ³⁰National Centre for Radio Astrophysics, Pune, India. ³¹These authors contributed equally: Tomas Cassanelli, Calvin Leung, Pranav Sanghavi. ³²Deceased: Jing Luo. ³³Unaffiliated: organisation name, city, and country. ✉ e-mail: calvin_leung@berkeley.edu



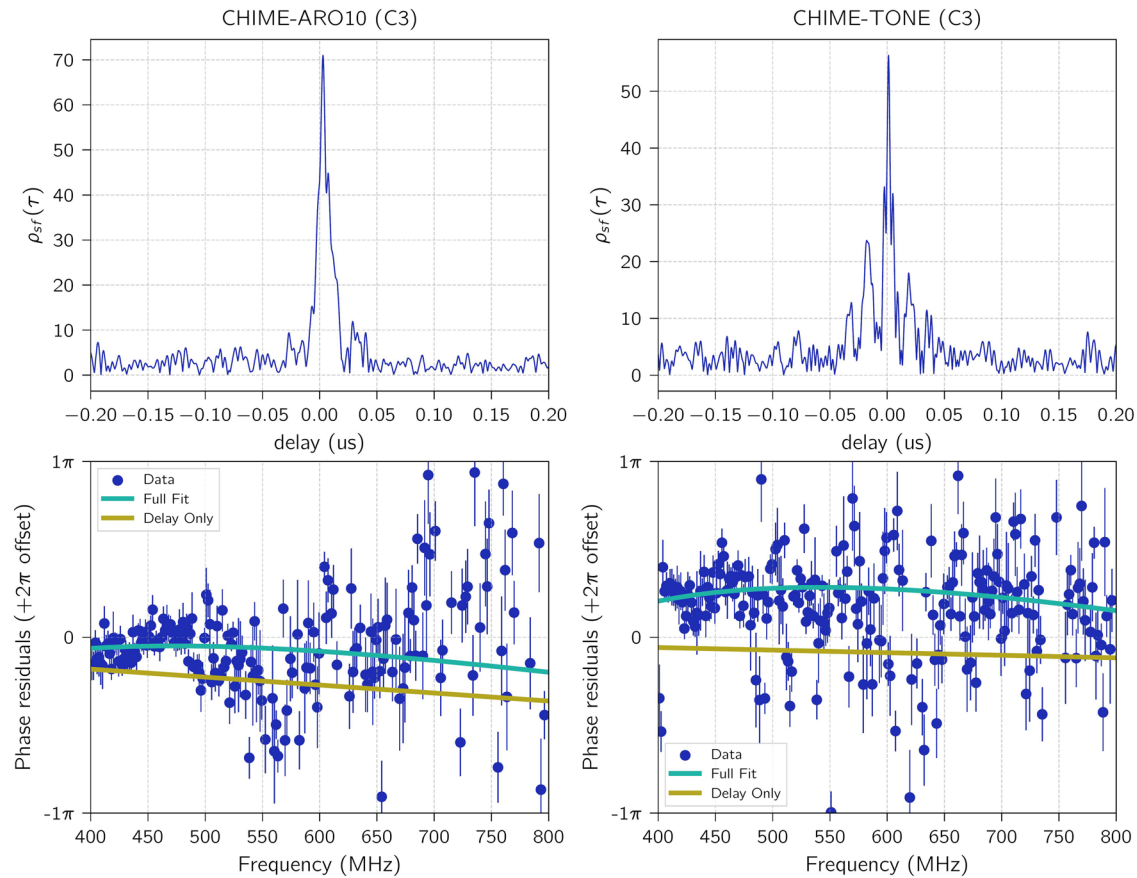
Extended Data Fig. 1 | Dynamic spectra of all observations. At each VLBI station we recorded five single pulses (including the FRB): Crab GPs which we refer to as C1–C4 in the several days surrounding FRB 20210603A. Each row corresponds to a different VLBI station (CHIME at the Dominion Radio Astrophysical Observatory, ARO10 at the Algonquin Radio Observatory, and TONE at the Green Bank Observatory). Timestamps show site-local clocks aligned to within $2.56 \mu\text{s}$ at a reference frequency of 800.0 MHz. Though the FRB is too faint to be detected at the testbeds alone, it is robustly detected in cross-correlation with CHIME at both stations. The intensity was adjusted by normalizing its standard deviation and setting the colour scale limits to the 1 and 99 percentile values of the data. Waterfall plots are shown downsampled to a

frequency resolution of 390.625 kHz and a time resolution of $25.6 \mu\text{s}$. The noisy radio frequency interference (RFI) channels in 700–750 MHz correspond to the cellular communications bands and the RFI channels at ≈ 600 MHz frequencies correspond to television transmission bands. These RFI channels were removed in our analysis and are highlighted with red strikes to the left of each waterfall plot. Symbols next to the telescope label in each waterfall plot indicate what each Crab pulse was used for. We use C2 on all baselines as a phase/delay calibrator, and C1 and C4 as rate calibrators for the CHIME–ARO10 and CHIME–TONE baselines respectively. We localized C3 as an end-to-end cross-check of our calibration solutions.




Extended Data Fig. 2 | Delay residuals measured from the CHIME-ARO10 and CHIME-TONE baselines. The graph shows the empirical uncertainty obtained by analysing earlier data sets^{3,4}, with CHIME-ARO10 data shown in the top row and CHIME-TONE data showed in the bottom row. Each point corresponds to

the residual delay after applying delay and phase corrections (CHIME-ARO10 is calibrated to 2020-10-22, and TONE is calibrated to 2021-02-18). The extracted delays have all been compensated for clock errors and for a clock rate error on the CHIME-ARO10 baseline.




Extended Data Fig. 3 | Calibrated visibilities from the Crab pulsar giant pulse (C3) used to validate our calibration solutions. We plot visibilities from the CHIME-ARO10 (left) and CHIME-TONE (right) baselines respectively. In each top panel, we plot the absolute value of the Fourier transform of the visibilities (that is the time-lag cross-correlation function $\rho(\tau)$) as a function of the delay referenced to the correlator pointing center). This shows a detection S/N exceeding 50 on each baseline. In each bottom panel we plot the phase of the calibrated visibilities $V[i, k]$, binned to 1.6 MHz resolution, with 1σ phase errors estimated from off-pulse scans ($N=10$) plotted as $\sigma[i, k]/V[i, k]$ (blue

points). In the bottom panels we overlay the phase model (Eq. (4)) evaluated at the parameters which maximize $L\Phi$, where we have fit for the ionosphere and the positions simultaneously (green “full fit” curve), as well as the phase model evaluated at the L_r position at zero ionosphere (yellow “delay only” curve). Since our correlator pointing is the L_r position, we would then expect the yellow “delay only” curve to be flat; note that our plotting code automatically unwraps all of the phases in each bottom panel by some amount automatically chosen to reduce phase wrapping, explaining the very small deviation from zero delay.



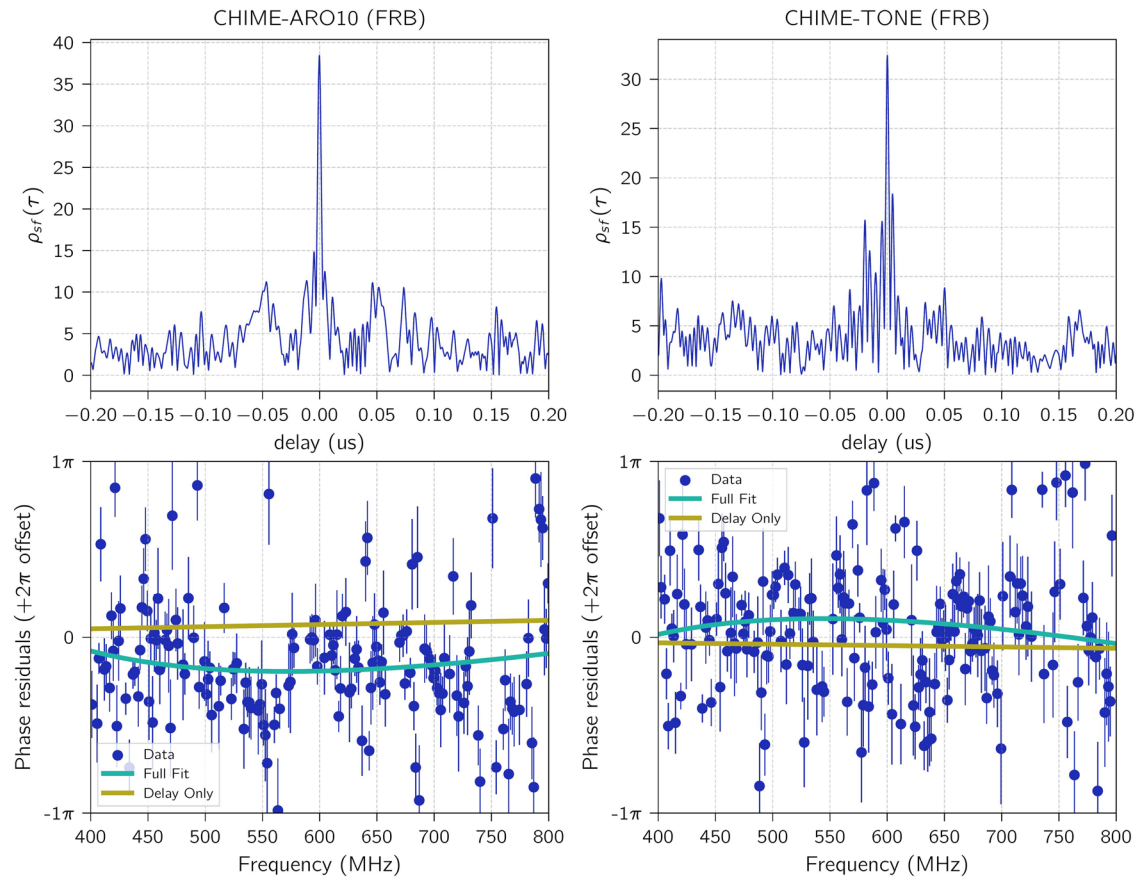
Extended Data Fig. 4 | The localization posterior of the Crab pulse (C3) as a function of R.A., Dec, and ΔDM_{CA} , and ΔDM_{CT} . Due to the extremely sparse sampling of the uv-plane, we bypass traditional methods of VLBI imaging, and directly fit the visibilities $V[i, k]$. Owing to our wide bandwidth, we see that the ionosphere parameters ΔDM are well-constrained even in the absence of external

information (for example, TEC maps or ionosphere priors). In the same spirit as a MCMC corner plot, each 2D plot shows the posterior marginalized over all except two axes. Calling these projections P , we colour evenly-spaced contours between $\log P = 0$ (the maximum value of each P is normalized to 1) and $\log P = -16$.



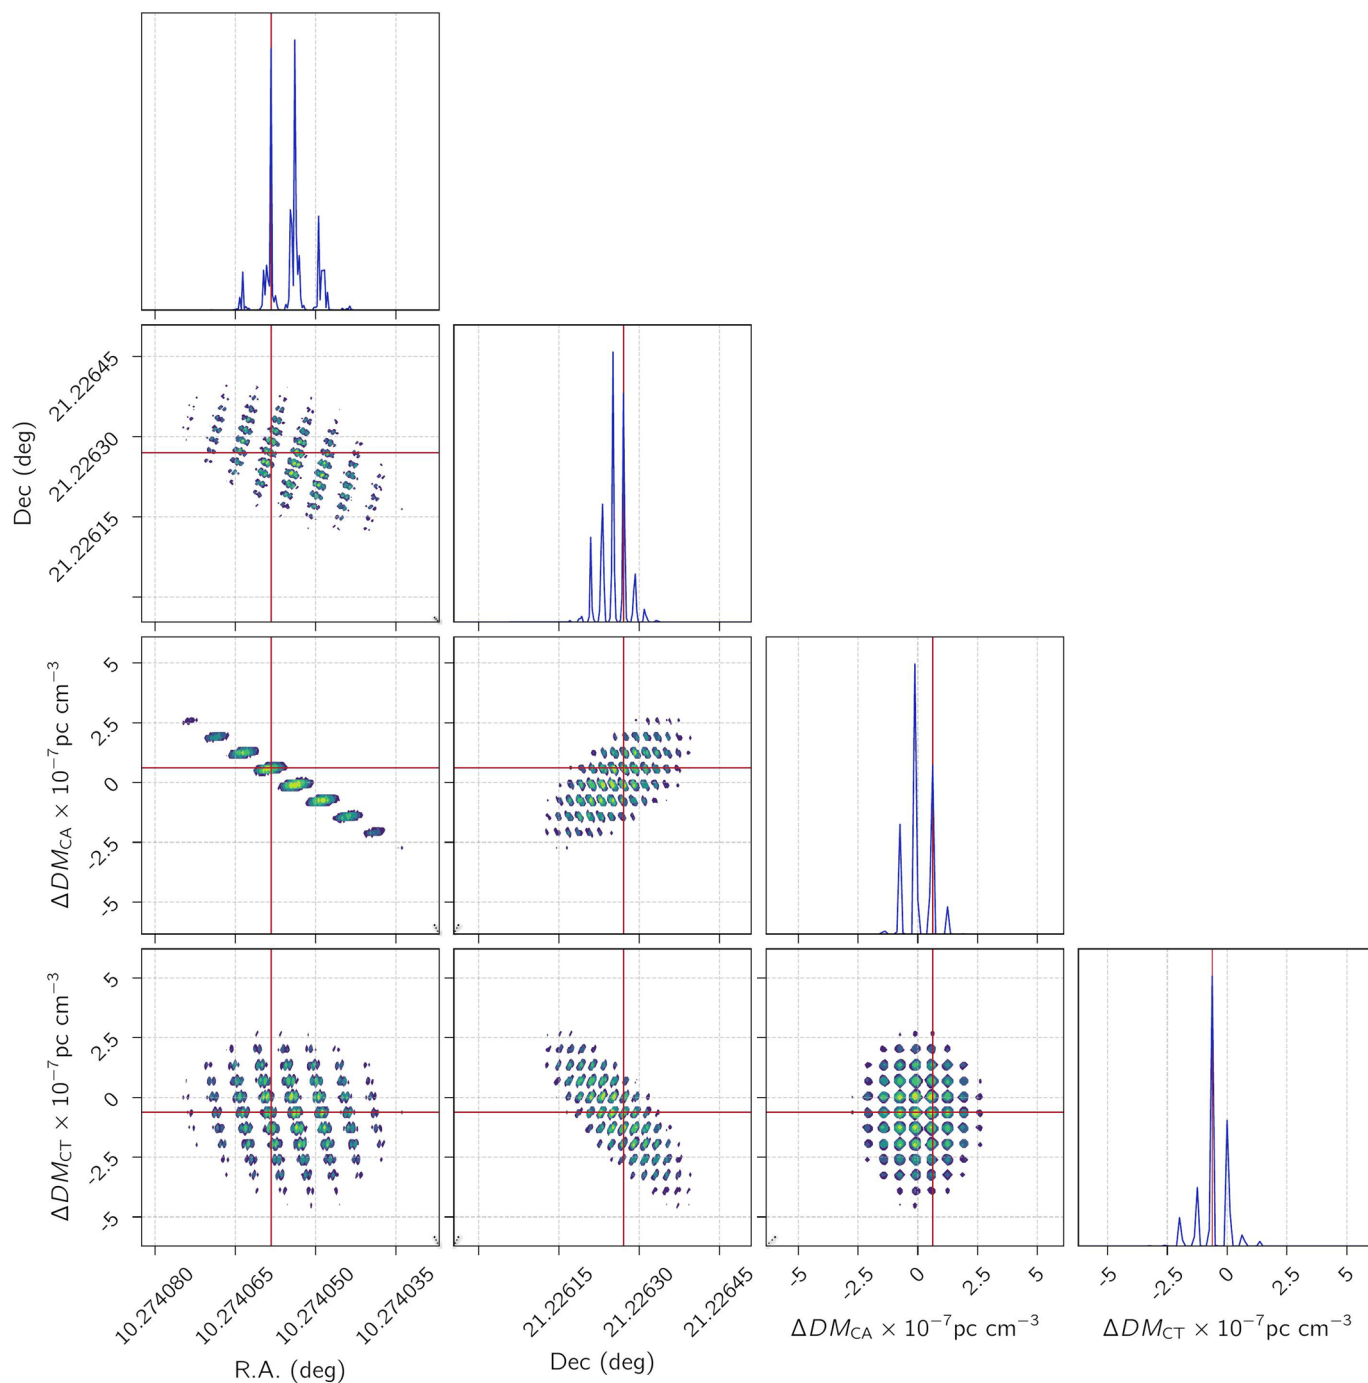
Extended Data Fig. 5 | Localization of C3 as an independent, end-to-end cross check of our VLBI calibration solution used to localize the FRB. Due to the extremely sparse sampling of the uv-plane, we avoid traditional imaging. We compare two localization methods: a delay-space χ^2 -minimization of the

residual delays left after calibration (+), and a visibility-space fitting of the phases (x). Both methods agree to within the true position of the Crab (star) within systematic uncertainties (ellipses).




Extended Data Fig. 6 | Calibrated VLBI fringes on FRB 20210603A from the CHIME- ARO10 and CHIME-TONE baselines respectively. We plot visibilities from the CHIME- ARO10 (left) and CHIME-TONE (right) baselines respectively. In each top panel, we plot the absolute value of the Fourier transform of the visibilities (that is the time-lag cross-correlation function $\rho(\tau)$) as a function of the delay referenced to the correlator pointing center). This shows a detection S/N exceeding 50 on each baseline. In each bottom panel we plot the phase of the calibrated visibilities $V[i, k]$, binned to 1.6 MHz resolution, with 1σ phase errors estimated from off-pulse scans ($N=10$) plotted as $\sigma[i, k]/V[i, k]$ (blue

points). In the bottom panels we overlay the phase model (Eq. (4)) evaluated at the parameters which maximize $L\Phi$, where we have fit for the ionosphere and the positions simultaneously (green “full fit” curve), as well as the phase model evaluated at the L_r position at zero ionosphere (yellow “delay only” curve). Since our correlator pointing is the L_r position, we would then expect the yellow “delay only” curve to be flat; note that our plotting code automatically unwraps all of the phases in each bottom panel by some amount automatically chosen to reduce phase wrapping, explaining the very small deviation from zero delay.



Extended Data Fig. 7 | The posterior localization contour of FRB 20210603A as a function of RA, Dec, and ΔDM_{CA} , and ΔDM_{CT} . The ionosphere parameters ΔDM are well- constrained even in the absence of external information (for example, TEC maps or ionosphere priors). In the same spirit as a MCMC corner

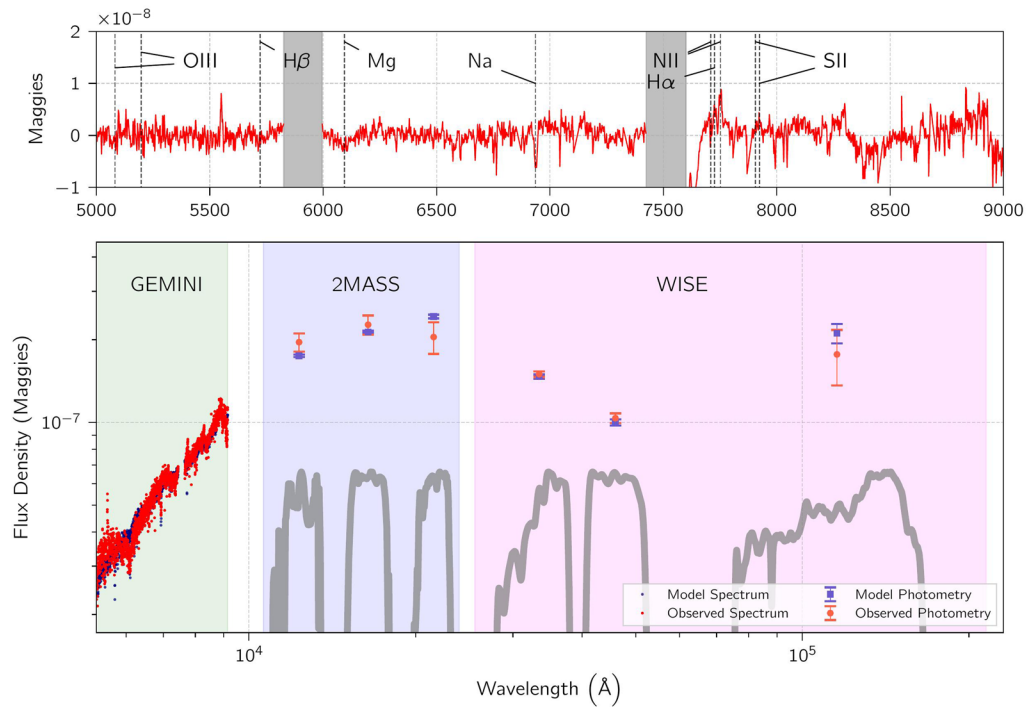
plot, each 2D plot shows the posterior marginalized over all except two axes. Calling these projections P , we colour evenly-spaced contours between $\log P = 0$ (the maximum value of each P is normalized to 1) and $\log P = -16$.



Extended Data Fig. 8 | Spatially resolved spectroscopy of the host galaxy.

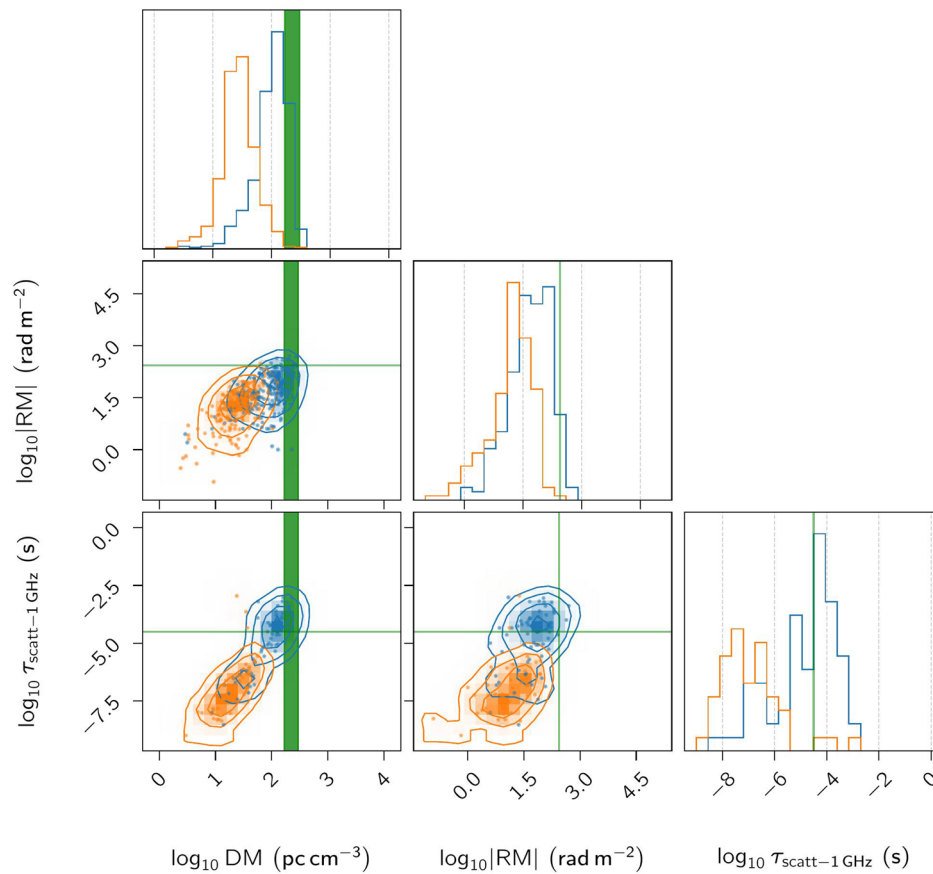
Optical image and spatially-resolved spectra of the host galaxy of FRB 20210603A acquired using CFHT MegaCAM and Gemini long-slit spectroscopy respectively. Pixel intensities are scaled linearly and normalized to reduce the saturation evident in Fig. 3. All spectra are given offsets in increments of $10^{-17} \text{ erg s}^{-1} \text{ cm}^{-2} \text{ \AA}^{-1}$. One spectrum is extracted from the bulge of the galaxy (spectrum b, centered at 0). There are additional eleven spectra extracted from the FRB side of

the galaxy (shown as positive offsets), and from the opposite side of the galaxy (shown as negative offsets), with offsets from the center of the galaxy in increments of 1 arcsec. All spectra are extracted using an aperture size of $1.5 \text{ arcsec} \times 1 \text{ arcsec}$, as represented on the galaxy image. Spectrum a is extracted within the vicinity of the FRB and represented by the shaded box a in the galaxy image. The twelve spectra and Gaussian fits to the $\text{H}\alpha$ and one of the NII emission lines, are plotted here after correcting for Milky-Way extinction.



Extended Data Fig. 9 | Spectral energy distribution of host galaxy. Gemini long-slit spectrum, integrated over the galaxy, with archival infrared photometry from 2MASS and WISE, plotted after correcting for extinction due to the host galaxy's inclination angle. Plotted alongside the observations (red) are the

best-fit model (blue) from Prospector, and the relative passbands for the 2MASS J, H, and Ks and WISE W1-W3 filters. Flux uncertainties are plotted by converting 1σ photometric errors reported by each catalogue.



Extended Data Fig. 10 | A visualization of propagation effects due to the Milky Way's disk, as measured via the ATNF Pulsar Catalogue. We plot joint distributions of DM, $|RM|$ and τ_{scatt} for Galactic pulsars for two different latitude ranges: $4^\circ \leq |b| \leq 10^\circ$ (blue) and $|b| \geq 20^\circ$ (orange) taken from the ATNF Pulsar Catalogue⁷⁹. Contour lines indicate 1, 2 and 3σ regions of this parameter space. Green regions/lines indicate estimates of equivalent quantities determined for

the host galaxy of FRB 20210603A, namely: DMhost, $|RM_{\text{host}}|$ and our upper limit on τ_{scatt} . DMhost, $|RM_{\text{host}}|$ and τ_{scatt} estimates are in the source frame with τ_{scatt} referenced at 1 GHz assuming a $\tau_{\text{scatt}} \propto \nu^{-4.4}$ relation used by ATNF. This shows that the burst properties of FRB 20210603A (DMhost, $|RM_{\text{host}}|$ and τ_{scatt} at 1 GHz), once corrected for extragalactic contributions, are similar to that of low-latitude ($4^\circ \leq |b| \leq 10^\circ$) Galactic pulsars.